Assessing the Aggregate Price Effects of Trade Using the Food Engel Curve*

Farid Farrokhi
Boston College

David Jinkins
Copenhagen Business School

Chong Xiang
Purdue University

February 27, 2026

Abstract

This paper constructs aggregate price indexes based on trade data and theory. We assess these Open-Economy Price Indexes (OPIs) using food Engel curves from an Almost Ideal Demand System. Estimating this demand system with household-level consumption data from the US, bias of any particular price index can be measured by deviation from Engel curves. Comparing OPIs with official CPI statistics, we find that the OPIs that use disaggregated import price data outperform CPI in our period of study.

Keywords: Food Engel Curves, Price Indices, Household-level Consumption

1 Introduction

The impact of international trade on aggregate prices is fundamental to trade theory. A substantial literature has developed methods to predict the aggregate price effects of trade, using detailed trade data (Feenstra, 1994; Broda and Weinstein, 2006; Redding and Weinstein, 2020) and calibrated quantitative general equilibrium models (Eaton and Kortum, 2002; Anderson and Van Wincoop, 2003; Arkolakis et al., 2012; Caliendo and Parro, 2015). In this paper, we build on this literature to construct open-economy price indices (OPIs) for final consumption. We then test whether the aggregate price changes predicted by canonical trade models are consistent with observed household consumption patterns, allowing us to compare the bias in OPIs to that of the conventional consumer price index (CPI)—which is subject to issues such as quality adjustment and the introduction of new products.

Since quantitative trade theories are designed to yield price indices in line with the gravity equation of trade, we require additional, equally well-established empirical regularities to assess their aggregate price predictions. To this end, we use the food Engel curve, also known as Engel's

*This research was supported by the Danish Independent Research Fund grant no. 8019-00031B. We thank Sunham Kim and Danyang Zhang for their excellent research assistance. We are grateful to Seungki Hong and Hamid Firooz as well as participants in various seminars for their helpful feedback. All errors are our own.

Corresponding author: David Jinkins, dj.eco@cbs.dk.

law (Engel, 1895)[1] indicating that increases in income lead to a lower share of expenditure on food. A body of literature has leveraged deviations from the food Engel curve to estimate bias in CPI (Costa, 2001; Hamilton, 2001), identify bias in purchasing power parity deflators across countries (Almås, 2012), and evaluate the accuracy of official statistics on inflation and growth (Nakamura et al., 2016). Like this literature, we estimate food Engel curves derived from the AIDS demand system, using U.S. household-level consumption data and measured price indices for food and non-food goods, based on either the CPI or OPIs. Because the estimated position of the Engel curve across household income levels depends critically on measured prices, year-to-year deviations from these curves reveal their bias relative to the unbiased "true" prices. Intuitively, the test is based on the idea that, if price indices are measured correctly, the relationship between food expenditure shares and real income should remain stable over time, so systematic shifts in that relationship suggest bias in the measured price index.

To incorporate trade theory and data into our analysis, we specify household demand using a nested structure. At the upper tier, households allocate spending between food and non-food composites, following the AIDS specification described above. Within the food and non-food categories, we assume further nested groupings that are homothetic, in line with the trade models we employ. These broader categories aggregate over more detailed industries, each of which combines domestic and international varieties within its corresponding industry.

We distinguish between trade models in constructing our price indexes by splitting our OPIs under two umbrellas which are theoretically equivalent but differ in their data requirements. The first, which we call OPI-D, emphasizes its greater reliance on domestic data. This approach uses observed changes in domestic prices and infers changes in import prices based on expenditure shares for foreign goods. The second approach, referred to as OPI-M, uses disaggregated data on unit values and quantities of imported products and infers changes in domestic prices from changes in the domestic share of expenditure.

We further distinguish two approaches to the calculation of OPI-Ms. Our first OPI-M specification follows the methodology of Feenstra (1994) and Broda and Weinstein (2006), which is an extension of the Sato-Vartia index (Vartia, 1976). We refer to the resulting index as OPI-M1. Our second specification, OPI-M2, follows Redding and Weinstein (2020)'s (RW) corrections of the earlier methodology. The Feenstra-Sato-Vartia methodology assumes away movements of unobserved demand shifters when assigning weights to the price of each product, whereas RW incorporates them in the calculation of price indices. By estimating the average biases of OPI-M1 and -M2 from the food Engel curve, we study whether, and by how much, RW's corrections move the overall price index closer to the true price index that consumers use in making consumption decisions.

The data used to construct the OPI indices come from product-level trade data on unit values and quantities taken from the BACI-CEPII database, and sectoral production data from STAN. For estimating food Engel curves, we use micro-data from the nationally representative Panel Study of

[1]A large literature, *e.g.* Deaton and Muellbauer (1980, 1986); Banks et al. (1997); Fajgelbaum and Khandelwal (2016), has shown that the food Engel curve is an appropriate characterization of household consumption for low-, middle- and high-income countries.

Income Dynamics (PSID) household consumption surveys. Our estimation also incorporates regional variation by using additional price data from the historical regional price information for goods and services from the Council for Community and Economic Research, as well as restricted data on the regional location of households in the PSID. Our detailed data provide the necessary variation for our estimation equation with second-order terms of log prices, and help us identify all the AIDS parameters.

Before evaluating the OPIs, we conduct two checks on our method. First, we find that the U.S. CPI introduces an average upward bias of 0.7 log points per year. This result aligns with previous studies of the U.S. CPI. For example, Gordon (2006) and Berndt (2006) report annual biases in the range of 0.7 to 0.9 percent. Second, when we use constant prices for all years in our estimation, our finding indicates a sizable underestimation of inflation, and the average bias is more than three times as large as that of the official U.S. CPI. These checks suggest that our novel data and estimation approach yield sensible results.

Our first key finding is that both of OPI-M indices exhibit less bias in tracking U.S. household consumption decisions than the official CPI. For instance, the average bias of OPI-M1 is more than 30% lower than that of the official CPI. In contrast, OPI-D shows a similar average bias to that of the CPI. Since OPI-M and OPI-D are theoretically two sides of the same coin, the discrepancy in their performance stems from differences in the data requirement for their measurement. In particular, measuring the purely domestic price component needed for OPI-D is inherently challenging, whereas there is detailed data available for calculating the import price component required for OPI-M. Specifically, detailed import price data allow us to incorporate the extensive margin in the calculation of the import price index. Consumers benefit not only from price reductions on existing varieties, but also from access to a broader range of products. In line with this argument, our decomposition, using OPI-M2, shows that when the extensive margin is removed, the bias increases to a level similar to that of the official CPI.

Our second key result is that OPI-M2 exhibits notably lower bias than OPI-M1, by about 40%. Since both indices use the same data, this difference must arise from how they account for changes in consumer tastes, as changes in product quality that shift demand beyond changes in product-level prices. Thus, the improvement in OPI-M2 can be attributed to the internally consistent quality measure developed by RW, rather than the Feenstra-Sato-Vartia index used in OPI-M1. This more accurate measurement of quality changes may also explain why the average bias of OPI-M2 is less than half of the official CPI. Indeed, a substantial body of literature, as reviewed below, has shown that quality bias is a major reason the U.S. CPI tends to overstate cost-of-living inflation.

Lastly, we leverage our estimated parameters of the AIDS preferences under our best-performing measured price, OPI-M2, to assess the effects of a hypothetical 10% increase in import prices on cost of living, using 2015 as the baseline year while holding domestic prices and income levels constant. We find that OPI-M2 would rise by 0.9% for food and 0.8% for non-food, reflecting their different exposures to trade. Calculating the cost-of-living (COL) index by household income, we find that across households, the average increase in COL index would be 0.8%. We also find that the changes in the COL index are slightly regressive, as poorer households are more affected due to their higher food consumption and the greater increase in food prices compared to non-food prices.

Related Literature. Our study relates to several strands of research. First, a literature uses the food Engel curve to estimate the biases in a single price index, such as the U.S. CPI or PPP indices, *e.g.*, Costa (2001); Hamilton (2001); Almås (2012). Relative to this literature, we estimate a fully specified demand system. This allows us to explicitly recognize the COL index as household specific. We also focus on the biases of the OPI indices derived from trade theory. Relatedly, a large body of research has shown that quality and new-goods biases are important reasons why the U.S. CPI overstates true inflation, *e.g.*, Moulton (1996); Hausman (2003).[2] Our results suggest that the estimated CPI biases could be related to mismeasured gains from trade, which, when measured based on trade theory, can capture adjustments in quality as well as the entry and exit of product varieties.

In turn, the impact of trade on prices has been the focus of a substantial body of literature. One approach, following Feenstra (1994), utilizes detailed product-level trade data, while another demonstrates the usefulness of sufficient statistics in quantitative general equilibrium trade models, as reviewed by Costinot and Rodríguez-Clare (2014).[3] Using these tools, several studies provide evidence on specific cases. For example, to study the impact of Chinese imports on U.S. consumer prices Amiti et al. (2020) employ the import-price-index approach, while Bai and Stumpner (2019) use the sufficient-statistics approach. In our analysis, we provide evidence that households' consumption decisions, as revealed by the food Engel curve, is consistent with real consumption gains predicted by this literature. We also clarify which specific margins contribute to reducing biases in OPI measures. Complementing this literature, Jaravel and Sager (2025) use detailed price microdata to causally link increases in Chinese import penetration to lower U.S. consumer prices and domestic markups. Rather than studying the underlying mechanisms behind consumer price changes at the level of products, we use household expenditure data to evaluate aggregate price indexes constructed directly from import unit values, trade shares, and domestic production data.

Our work also relates to research on the exact cost-of-living index under non-homothetic preferences. Jaravel and Lashkari (2024) develop nonparametric methods to measure changes in consumer welfare that allow demand patterns to vary flexibly with income and other observed characteristics, and Baqaee et al. (2024) show how to recover money-metric utility from repeated cross sections. In both cases, the well-measured price series enter as observed inputs to the welfare/cost-of-living calculations, and the methods do not separately model or identify systematic price-measurement error from the same demand data. Atkin et al. (2020) use relative Engel curves and well-measured prices

[2]See also Shapiro and Wilcox (1996); Boskin et al. (1997); Moulton et al. (1997); Berndt (2006); Gordon (2006); Greenlees and McClelland (2011); Melser and Syed (2016).

[3]See also Arkolakis et al. (2008); Balistreri et al. (2011); Ossa (2015); Hsieh and Ossa (2016); Levchenko and Zhang (2016) and Giri et al. (2021), among others. Some studies have also examined the pro-competitive effects of trade through markups (*e.g.*, Feenstra and Weinstein (2017)), which we do not explore in our analysis. Others have investigated various channels beyond the canonical models we consider, such as search frictions (*e.g.*, Krolikowski and McCallum, 2021; Eaton et al., 2021, 2022), specifics of trade in commodities (*e.g.*, Farrokhi 2020; Fally and Sayre 2018), global value chains (*e.g.*, Antràs and De Gortari, 2020), and multinational production (*e.g.*, Arkolakis et al., 2018; Du and Wang, 2021), among other channels.

for a subset of goods to estimate cost-of-living indexes across the income distribution.[4] In sum, this literature shows it is difficult to address non-homothetic preferences and mismeasured prices simultaneously, and thus typically assumes that well-measured prices are available.[5] In comparison, we start from the premise that well-measured prices, adjusted for quality and variety, are difficult to obtain for most merchandise goods, and so our goal and approach differ from this literature. Still, our results and those of Atkin et al. (2020) both show that price indices correcting for quality and variety exhibit lower increases over time relative to the official CPI.

Lastly, to assess the validity of quantitative trade models, Kehoe et al. (2015) and Kehoe et al. (2017) compare observed changes in trade flows with those predicted by applied general equilibrium (AGE) models in response to specific trade policies like NAFTA. In turn, Adão et al. (2025) introduce a goodness-of-fit measure that accounts for the fact that observed data reflect not only the effects of a specific policy of interest but also changes in other policies and market conditions. Our work contributes to this literature by using deviations from the food Engel curve as a basis for evaluating trade models' predictions of price indices.

Roadmap. The remainder of the paper is organized as follows. Section 2 explains how we estimate the food Engel curve to assess and compare biases in measured price indices. Section 3 outlines our basic approach to calculating OPIs with a simple example. Section 4 details our full model and the OPI calculation procedure. Section 5 describes the data and presents preliminary analysis. Sections 6 presents our results and discuss their implications. Section 7 considers a few extensions to our analysis. Section 8 concludes.

2 Comparing Price Indices Using Household Consumption Data

Our approach is inspired by Costa (2001) and Hamilton (2001), but it also introduces some key differences by deriving our estimation equation explicitly from an AIDS demand system. We begin by outlining our method for estimating the food Engel curve, which allows us to recover the bias in measured prices. We then describe how we rank and compare these measured prices.

2.1 Households' Consumption and Food Engel Curve

Consider households indexed by h in region r and year t, with their preferences represented by an Almost Ideal Demand System (AIDS) à la Deaton and Muellbauer (1980). Let E_h^t and $\tilde{\mathbf{P}}_r^t = (\tilde{P}_{r,F}^t, \tilde{P}_{r,N}^t)$ denote household expenditure, and "true" price indices of food (F) and non-food (N), which vary by year t and region r. The indirect utility function for the corresponding household is

[4] In theory, computation remains valid without well-measured prices under an orthogonality condition, but well-measured prices are needed to test this condition. In the data, Atkin et al. (2020) show that this orthogonality condition is a good approximation for rural Indian households.

[5] See also Fajgelbaum and Khandelwal (2016); Borusyak and Jaravel (2018); Almås et al. (2018); Argente and Lee (2021); Auer et al. (2022).

given by:

$$U_h^t = \mathrm{U}(E_h^t, \tilde{\mathbf{P}}_r^t) = \frac{\ln E_h^t - \ln \Gamma(\tilde{\mathbf{P}}_r^t)}{\Lambda(\tilde{\mathbf{P}}_r^t)}, \tag{1}$$

where $\Gamma(.)$ and $\Lambda(.)$ are given by:

$$\begin{cases} \ln \Gamma(\tilde{\mathbf{P}}_r^t) = \alpha_0 + \alpha_F \ln \tilde{P}_{r,F}^t + (1-\alpha_F) \ln \tilde{P}_{r,N}^t + \frac{\gamma}{2}\left(\ln \tilde{P}_{r,F}^t - \ln \tilde{P}_{r,N}^t\right)^2 \\ \ln \Lambda(\tilde{\mathbf{P}}_r^t) = \ln \beta_0 + \beta_F \ln \tilde{P}_{r,F}^t - \beta_F \ln \tilde{P}_{r,N}^t \end{cases} \tag{2}$$

Here, $\{\alpha_0, \alpha_F, \gamma, \beta_0, \beta_F\}$ are time-invariant parameters, $\Gamma(.)$ represents a reference for the costs of subsistence, and $\Lambda(.)$ is a price aggregator that scales the level of utility. By utility maximization, the expenditure share of food, $\lambda_{h,F}^t \equiv \tilde{P}_{r,F}^t Q_{h,F}^t / E_h^t$, equals:

$$\lambda_{h,F}^t = \delta_0 + \beta_F \ln E_h^t + \delta_F \ln \tilde{P}_{r,F}^t + \delta_N \ln \tilde{P}_{r,N}^t + \delta_X \left(\ln \tilde{P}_{r,F}^t - \ln \tilde{P}_{r,N}^t\right)^2, \tag{3}$$

where the auxiliary coefficients $\{\delta_0, \delta_F, \delta_N, \delta_X\}$ are defined as:

$$\delta_0 \equiv \alpha_F - \beta_F \alpha_0, \quad \delta_F \equiv \gamma - \beta_F \alpha_F, \quad \delta_N \equiv -\gamma - \beta_F (1-\alpha_F), \quad \delta_X \equiv -\frac{\beta_F \gamma}{2} \tag{4}$$

Equations (3)-(4) are the theoretical motivation for a large body of work that estimates food Engel curves. The sign and magnitude of the income coefficient, β_F, indicate whether food is a luxury ($\beta_F > 0$), a necessity ($-\overline{\lambda}_F < \beta_F < 0$, where $\overline{\lambda}_F$ is the mean food expenditure share), or an inferior good ($\beta_F < -\overline{\lambda}_F$). The coefficient, γ, can be either positive or negative, depending on the price elasticity of demand for food. Specifically, Equation (3) implies that *(i)* the income elasticity of food for the average household is:

$$\frac{\partial \ln \overline{Q}_F}{\partial \ln E} = 1 + \frac{\beta_F}{\overline{\lambda}_F}, \tag{5}$$

where $\overline{Q}_F$ is the mean food consumption quantity, and *(ii)* the own price elasticity of food for the average household is:

$$\frac{\partial \ln \overline{Q}_F}{\partial \ln \tilde{P}_F} = \frac{1}{\overline{\lambda}_F}\left(\gamma - \beta_F \left(\alpha_F + \ln \frac{\tilde{P}_F}{\tilde{P}_N}\right)\right) - 1. \tag{6}$$

2.2 Evaluating Biases in Measured Prices Using Food Engel Curves

To estimate the food Engel curve, Equation (3), we use data on observed households' total expenditure, E_h^t, food expenditure share, $\lambda_{h,F}^t$, and "measured" prices of food and non-food, denoted by $\mathbf{P}_r^t = (P_{r,F}^t, P_{r,N}^t)$. The challenge is that any vector of measured prices is constructed subject to potential biases relative to the vector of true prices, $\tilde{\mathbf{P}}_r^t = (\tilde{P}_{r,F}^t, \tilde{P}_{r,N}^t)$. Specifically, the log true price can be expressed as the sum of log measured price and a potential bias:

$$\ln \tilde{P}_{r,j}^t = \ln P_{r,j}^t + \ln B_{r,j}^t, \qquad j = F, N; \tag{7}$$

where $B^t_{r,j}$ denotes the bias, which can be further decomposed to regional and annual components as:

$$\ln B^t_{r,j} = b_{r,j} + b^t_j + \varepsilon^t_{r,j} \qquad j = F, N. \tag{8}$$

Inserting Equations (7)-(8) into Equation (3) delivers food Engel curves evaluated at measured rather than true prices. To estimate the resulting equation, we make the following assumption:

Assumption 1. *(a) Region- and year-specific components of biases in Equation (8) are common between measured prices of food and non-food, i.e., $b_{r,j} = b_r$, $b^t_j = b^t$ for $j = F, N$. (b) The error terms $\varepsilon^t_{r,F}$ and $\varepsilon^t_{r,N}$ and their squared differences are uncorrelated with income, measured prices and region or year fixed effects.*

Assumption 1 underpins the identification of food Engel curves in our estimation. Part *(a)* states that biases over time and across regions are *common* across goods. Intuitively, we have only one Engel curve at each region and time, so we cannot identify more than one region and one time fixed effect. Part *(b)* allows us to consistently estimate the coefficients on income and prices.

Under Assumption 1, the food Engel equation, (3), evaluated at measured prices, becomes:

$$\lambda^t_{h,F} = \delta_r + \delta^t + \beta_F \left(\ln E^t_h - \ln P^t_{r,N}\right) + \delta_F \left(\ln P^t_{r,F} - \ln P^t_{r,N}\right) + \delta_X \left(\ln P^t_{r,F} - \ln P^t_{r,N}\right)^2 + u^t_h, \tag{9}$$

where u^t_h is mean zero and uncorrelated with the regressors, the region fixed effect is $\delta_r \equiv -\beta_F b_r + v_r$ and the year fixed effect, which is referred to in the literature as "drift", corresponds to $\delta_t \equiv -\beta_F b^t$. Note that δ^t is normalized to 0 for the initial year in the sample, $t = 0$, with the aid of regional fixed effects. Estimating (9) pins down $\hat{\beta}_F$ and $\hat{\delta}^t$ which we use to recover the common annual bias of prices:

$$\hat{b}^t = \frac{1}{\hat{\beta}_F}\hat{\delta}^t \tag{10}$$

Equations (9) and (10) allow us to rank and compare across different vectors of measured prices. To be specific, let the subscript m denote a specific vector of measured prices, such as official CPI, or an open-economy index, labeled as "OPI", constructed based on the gains-from-trade analyses. We compute the vector of food and non-food prices, $(P^t_{r,F,m}, P^t_{r,N,m})$, use them to estimate Equation (9), recover the biases, $\hat{b}^t_m$, according to Equation (10), and compute the root-mean square of the biases based on the following expression:

$$\text{RMSB}_m = \left(\frac{1}{T}\sum_{t=1}^{T}(\hat{b}^t_m)^2\right)^{\frac{1}{2}}, \tag{11}$$

where T is the number of years in the sample. Our logic is simple: if an OPI yields a sufficiently small RMSB_m, it is empirical evidence that the underlying gains-from-trade framework that generates this OPI is consistent with U.S. household consumption decisions. Needless to say, whether RMSB_m is "large" or "small" is relative, and so we use the official U.S. CPI as the benchmark; *i.e.* we evaluate each OPI against CPI using the metric of (11). We will also supplement the numerical values of RMSB_m with plots of $\hat{b}^t_m$ and their confidence intervals.

Lastly, we can recover γ and α_F using Equation (4) and estimates of $\hat{\delta}_F$ and $\hat{\delta}_X$,

$$\hat{\gamma} = -\frac{1}{\hat{\beta}_F} \times 2\hat{\delta}_X, \qquad \hat{\alpha}_F = -\frac{1}{\hat{\beta}_F} \times \left(\hat{\delta}_F + \frac{2\hat{\delta}_X}{\hat{\beta}_F} \right) \tag{12}$$

Discussion We now explicitly compare our approach with the literature, primarily Costa (2001), Hamilton (2001), and Nakamura et al. (2016), among others who have relied on food Engel curves for empirical study of price indices. To start, we re-write Equation (3) as:

$$\lambda_{h,F}^t = \alpha_F + \beta_F \left(\ln E_h^t - \ln \Gamma(\tilde{\mathbf{P}}_r^t) \right) + \gamma \left(\ln \tilde{P}_{r,F}^t - \ln \tilde{P}_{r,N}^t \right), \tag{13}$$

where $\ln \Gamma(\tilde{\mathbf{P}}_r^t)$, given by Equation (2), serves as an income deflator. Using the first-order approximation $\ln \Gamma(\tilde{\mathbf{P}}_r^t) \approx \ln \tilde{P}_r^t = \alpha_0 + \alpha_F \ln \tilde{P}_{r,F}^t + (1-\alpha_F) \ln \tilde{P}_{r,N}^t$, and by specifying the bias as in Equation (7) and adopting identification assumptions similar to Assumption 1, the literature reaches to an estimation equation of the following form:

$$\lambda_{h,F}^t = \delta_r + \delta^t + \beta_F \left(\ln E_h^t - \ln P^t \right) + \delta_F \left(\ln P_{r,F}^t - \ln P_{r,N}^t \right) + u_h^t. \tag{14}$$

The differences between our approach and that of the literature are as follows. First, our specification, Equation (9), does not make the approximation for $\ln \Gamma(\tilde{\mathbf{P}}_r^t)$, and so explicitly includes the square of the log of the relative price of food to non-food. This allows us to recover the bias in measured price indices via Equation (10), as the literature does, as well as to retrieve the values of the underlying parameters of the AIDS preferences of (1)-(2), via Equation (12). These parameter values are important for our welfare calculations and counterfactuals in sub-section 6.2 below.

Second, the literature uses the recovered bias from (10) to construct the single price index of $P^t + \hat{b}^t$, and interprets it as the true price index of the economy. In comparison, we do not aggregate the prices of food and non-food, because we recognize that AIDS preferences imply that the cost-of-living index varies with household income.[6] Instead, we use the recovered bias from (10) to evaluate how a vector of measured price indices of food and non-food compares with another vector, as revealed by the Food Engel curve.

In summary, relative to the literature, our empirical approach has a tighter connection with the underlying preferences and is more akin to structural estimation. Later, in sub-section 5.2, we clarify that we also use disaggregated data on the prices of food and non-food for our estimation.

3 Sketching the Idea: Open-Economy Price Indices

In this section, we use a toy model to illustrate how we construct the price indices that are implied by the gains-from-trade literature. We leave out the time superscript to keep our exposition compact. We set up our full model in the next section.

[6] This can be seen from Equation (1) wherein the change to utility, $\mathrm{d} \ln U_h^t$, differs from the change to income E_h^t deflated by $\Gamma(\tilde{\mathbf{P}}_r^t)$, *i.e*, $\mathrm{d}(\ln(E_h^t - \ln \Gamma(\tilde{\mathbf{P}}_r^t))$.

Suppose that the world economy consists of multiple countries, and a single sector. Goods are differentiated by country of production, indexed by $j = 1, ..., N$. Preferences are CES with substitution elasticity σ and demand shifters b_j. Consider a country called Home. Then, the CES price index in Home is

$$P = \left[\sum_{j=1}^{N} b_j \left(P_j\right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} . \tag{15}$$

A large body of work in the gains-from-trade literature goes after the same objective, which is, in its simplest form, the change to the CES price index specified in equation (15). We denote this change between two time periods by $\hat{P} = P^1/P^0$, which we refer to as OPI. We now outline two approaches to compute OPI.

3.1 OPI-D

Our first approach draws on the literature that uses sufficient statistics à la Arkolakis et al. (2012) to compute general equilibrium changes to prices. Specifically, the change to the price index, $\hat{P}$, can be expressed as:

$$\hat{P} = \hat{P}^D \left(\hat{\pi}^D\right)^{\frac{1}{\sigma-1}} , \qquad \text{[OPI-D]} \tag{16}$$

This equation shows that, given data on the change in domestic expenditure share $\hat{\pi}^D$ and the price of the domestic variety $\hat{P}^D$, we can compute the change in the overall price index, $\hat{P}$. Intuitively, when the domestic expenditure share decreases, it suggests that the price index for imported varieties has declined relative to the price of the domestic variety. This relationship is governed by the elasticity parameter $(\sigma - 1)$, which is why $\hat{P}^D$ must be adjusted by the term $\left(\hat{\pi}^D\right)^{\frac{1}{\sigma-1}}$. We refer to the OPI based on Equation (16) as OPI-D, where "D" stands for domestic.

A special case of Equation (16) delivers the gains from trade, defined as the loss in real income from moving the economy to autarky. This case, referred to in the literature as the ACR formula, corresponds to $\hat{\pi}^D = 1/\pi^D$ where π^D is the baseline (observed) value of the domestic expenditure share, and $\hat{P}^D = \hat{w}$ where $\hat{w}$ is the change to the nominal wage in Home. In that case, the change to real income, $\hat{w}/\hat{P}$, is given by $\left(\pi^D\right)^{\frac{1}{\sigma-1}}$.

3.2 OPI-M

Our second approach draws on the literature that uses detailed import data to examine the gains from trade. Here, we use data on foreign variety-level prices, p_j, and demand shifters, b_j, recovered from import shares, to compute the OPI. We refer to the resulting price index as OPI-M, where "M" stands for imports.

Specifically, varieties available to Home can be partitioned into domestic and imported varieties, whose price indices are, respectively, P^D and $P^M \equiv \left[\sum_{j \neq D} b_j \left(p_j\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$. This allows us to re-write Equation (15) as

$$\hat{P} = \hat{P}^M \left(\hat{\pi}^M\right)^{\frac{1}{\sigma-1}} , \qquad \text{[OPI-M]} \tag{17}$$

where $\pi^M = 1 - \pi^D$ is the imported share of expenditure. To see the intuition, consider a decrease in the import share, π^M, as households shift to domestic goods. This suggests domestic prices have fallen relative to import prices. Thus, we adjust the import price change downward by the import share decrease to accurately reflect the overall index change, OPI-M.

Equation (17) is the mirror image of Equation (16). This result is intuitive: OPI-M and OPI-D are two sides of the same coin. Although both calculate the change in the CES price index in Equation (15), they use different data. As a result, in practice they will likely generate different values of OPI, and so provide two complementary approaches to taking the gains-from-trade predictions to household consumption data.

4 A Framework for Computing Open-Economy Price Indices

Our full model is designed to bring the simple model of Section 3 to data. To this end, we begin with specifying the space of products. This specification is motivated by (i) the availability of micro data that can be used to construct the OPIs, and (ii) the manner in which non-homothetic preferences, which motivate the food Engel curve, can be combined with homothetic CES preferences.

Figure 1: Product Space

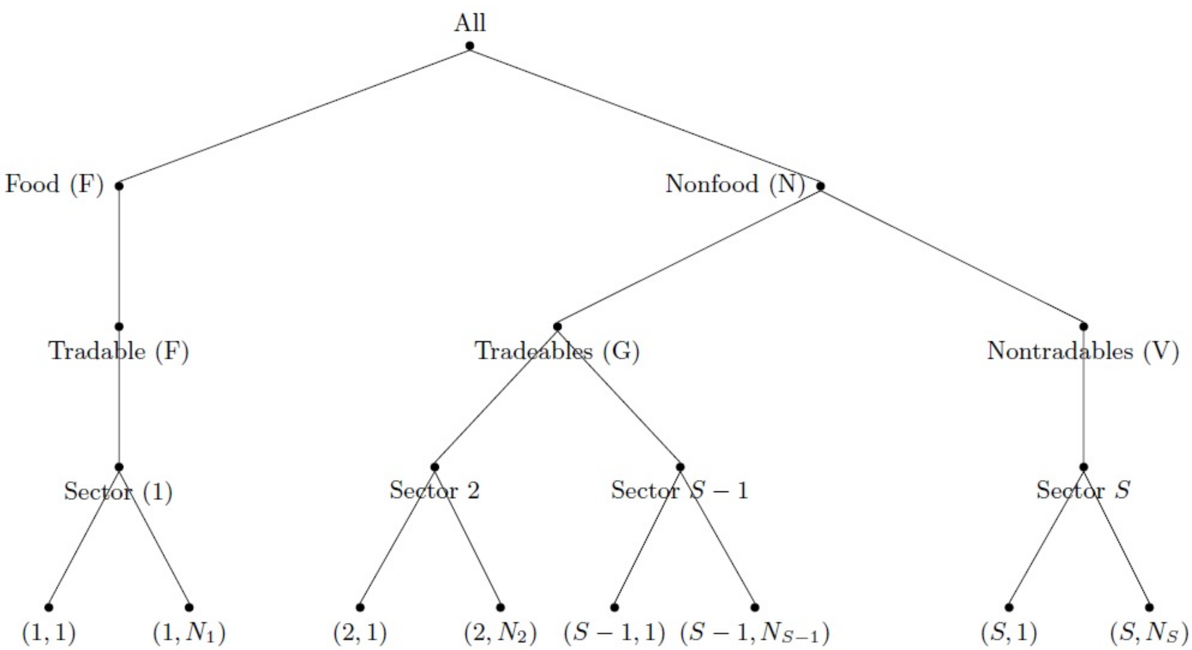

Notes: This figure shows the space of products and the nesting structure used in our analysis. N_1, N_2 etc. denote the numbers of goods within each sector.

We illustrate the space of products in Figure 1. The consumption of each household at time t consists of two broad groups, food (F), and non-food (N). This non-homothetic upper-tier demand gives rise to the food Engel curve as discussed in Section 2.1. The non-food group includes two categories, tradeable merchandise (G) and nontradeable services (V), and the food group is its own category. We model services as non-tradable because the data for import prices are for merchandise only. Within each category $c \in \{F, G, V\}$, there are multiple sectors indexed by s, with the food category F consisting of sector $s = \{1\}$, non-food tradeable category consisting of $s = \{2, ..., S-1\}$, and nontradeable non-food V consisting of $s = \{S\}$.

We next explain how we follow the trade literature to construct different specifications of the OPI-D and OPI-M indices for the homothetic categories F, G and V. We then show how we aggregate across categories to the food and non-food groups, the highest aggregation level that is not income specific, given the non-homothetic preferences between groups.

We assume that the sub-utility for sectors $s = \{1, ..., S-1\}$ is CES with substitution elasticity σ_s. We can then express the price index of sector s by adding the sector subscript s and time superscript t to equations (16) and (17), namely:

$$\hat{P}_s^t = \begin{cases} \hat{P}_s^{D,t} (\hat{\pi}_s^{D,t})^{\frac{1}{\sigma_s - 1}} & \text{[OPI-D]} \\ \hat{P}_s^{M,t} (\hat{\pi}_s^{M,t})^{\frac{1}{\sigma_s - 1}} & \text{[OPI-M]} \end{cases} \tag{18}$$

To calculate OPI-M and OPI-D indices of $\hat{P}_s^t$, we require data on changes to imported and domestic expenditure shares, $\pi_s^{M,t} = 1 - \pi_s^{D,t}$, and estimates of trade elasticities, σ_s, as well as sector-level import and domestic price indices, $P_s^{M,t}$ and $P_s^{D,t}$. Below, we present further assumptions that we invoke to calculate OPI-M and OPI-D at the level of sectors, and then spell out our aggregation. We discuss our sources of data and estimates in Section 5.

OPI-D. We calculate OPI-D according to the first line of equation (18). To this end, we require estimates of trade elasticity $(\sigma_s - 1)$ and observed changes to domestic expenditure shares $\hat{\pi}_s^{D,t}$, in addition to measures of sectoral domestic price, $\hat{P}_s^{D,t}$, as discussed in Section 5.

OPI-M. We consider two major approaches in the literature to calculate import price indices, $P_s^{M,t}$. The first follows Feenstra (1994) and Broda and Weinstein (2006) which we refer to as "FBW". The second follows Redding and Weinstein (2020) which we refer to as "RW".[7]

To compute $P_s^{M,t}$, we add a CES layer within each tradable sector. With this additional layer, $P_s^{M,t}$ captures the effects of variety entry and exit more accurately. To be specific, for the sectors in the tradable categories of food, F, and other merchandise, G, $s = \{1, ..., S-1\}$, $P_s^{M,t}$ is a CES aggregate across imported goods, indexed by g, whose price indices are $P_{gs}^{M,t}$. The imported goods are shown as the bottom-most nodes in Figure 1, and the substitution elasticity among them, within sector s, is σ_s^M. This implies that we can construct the sector-level price, $P_s^{M,t}$, from goods-level prices, $P_{gs}^{M,t}$, using the standard Sato-Vartia aggregation:

$$\hat{P}_s^{M,t} = \prod_{g \in \Omega_s} (\hat{P}_{gs}^{M,t})^{d_{gs}^t}, \qquad \text{where} \qquad d_{gs}^t = \frac{(\lambda_{gs}^t - \lambda_{gs}^0)/(\ln \lambda_{gs}^t - \ln \lambda_{gs}^0)}{\sum_{g \in \Omega_s} (\lambda_{gs}^t - \lambda_{gs}^0)/(\ln \lambda_{gs}^t - \ln \lambda_{gs}^0)} \tag{19}$$

In Equation (19), Ω_s is the set of goods within sector s which does not change over time. λ_{gs}^t

[7] For details on our formulas and derivations in this subsection, see Appendix B.1.

is the share of good g in import value within sector s in time t, and d^t_{gs} is the corresponding Sato-Vartia weights by good, computed as the log mean of λ^t_{gs} and λ^0_{gs}, the goods' import-value shares in time t and in the base period of 0. Equation (19) is common to both the FBW and RW procedures.

Each good, in turn, is differentiated into varieties, indexed by j, whose prices are $p^t_{j,gs}$. The substitution elasticity between the varieties within good g is σ_g, and their taste parameters are $b^t_{j,gs}$. These tastes can reflect a variety's quality, or any other demand shock. The set of varieties within g, $\Omega^{M,t}_{gs}$, may change over time. Let Ω^M_{gs} denote the common set, or the set of varieties that are present in both time t and the base period 0. The good-level import price index $\hat{P}^{M,t}_{gs}$ can be expressed as:

$$\hat{P}^{M,t}_{gs} = (\widehat{\lambda}^{*t}_{gs})^{\frac{1}{\sigma_g - 1}} \prod_{j \in \Omega^M_{gs}} (\widehat{p}^t_{j,gs})^{d^t_{j,gs}}, \qquad \text{where} \qquad \lambda^{*t}_{gs} = \frac{\sum_{j \in \Omega^M_{gs}} p^t_{j,gs} q^t_{j,gs}}{\sum_{j \in \Omega^{M,t}_{gs}} p^t_{j,gs} q^t_{j,gs}}. \tag{20}$$

In Equation (20), $p^t_{j,gs} q^t_{j,gs}$ is the import value of variety j at time t, and λ^{*t}_{gs} is the import-value share of the common set in t. The first factor on the right-hand side of Equation (20) captures the contribution to the good-level import price index from variety entry and exit. Intuitively, a rise in import values on new varieties at t implies a lower λ^{*t}_{gs} relative to λ^{*0}_{gs}, and leads to a fall in the overall index, $\hat{P}^{M,t}_{gs}$. This factor is the same for both the FBW and RW procedures. The second factor in (20) captures the price changes that come from continuing varieties and it is different for the FBW and RW procedures.

The FBW procedure applies the standard Sato-Vartia weights of log means to the variety level:

$$d^t_{j,gs} = \frac{(\lambda^t_{j,gs} - \lambda^0_{j,gs}) / (\ln \lambda^t_{j,gs} - \ln \lambda^0_{j,gs})}{\sum_{j \in \Omega^M_{gs}} (\lambda^t_{j,gs} - \lambda^0_{j,gs}) / (\ln \lambda^t_{j,gs} - \ln \lambda^0_{j,gs})}, \tag{21}$$

where $\lambda^t_{j,gs}$ is the import-value share of variety j within the common set, Ω^M_{gs}. Within this common set, the FBW procedure assumes, in addition, that the taste parameters of every variety remain unchanged over time; *i.e* $\hat{b}^t_{j,gs} = 1$ for all $j \in \Omega^M_{gs}$. Redding and Weinstein (2020) point out that the use of the weights from Equation (21) are incompatible with the assumption of constant taste parameters, and show, instead, that their changes over time can be recovered from the residual of import-value shares conditional on observed prices, under the weaker assumption that taste changes are zero, on geometric average, within the common set. To show this point explicitly, let $\widetilde{x}$ denote the simple geometric mean across varieties in the common set for variable x. Then the assumption of zero taste changes on average can be expressed as $\widetilde{b}^t_{gs} = \widetilde{b}^0_{gs}$ for all t, and the over-time changes to variety-level taste parameters can be recovered as follows:

$$\left(\ln b^t_{j,gs} - \ln b^0_{j,gs} \right) = \ln \left[\left(\frac{\lambda^t_{j,gs}}{\widetilde{\lambda}^t_{gs}} \right) \Big/ \left(\frac{\lambda^0_{j,gs}}{\widetilde{\lambda}^0_{gs}} \right) \right] - (1 - \sigma_g) \ln \left[\left(\frac{p^t_{j,gs}}{\widetilde{p}^t_{gs}} \right) \Big/ \left(\frac{p^0_{j,gs}}{\widetilde{p}^0_{gs}} \right) \right] \tag{22}$$

This expression implies the following modified Sato-Vartia weights for the RW procedure:

$$d^t_{j,gs} = \frac{\left(\lambda^t_{j,gs} - \lambda^0_{j,gs}\right) \Big/ \left((\ln \lambda^t_{j,gs} - \ln \lambda^0_{j,gs}) - (\ln b^t_{j,gs} - \ln b^0_{j,gs})\right)}{\sum_{j \in I^M_g} \left(\lambda^t_{j,gs} - \lambda^0_{j,gs}\right) \Big/ \left((\ln \lambda^t_{j,gs} - \ln \lambda^0_{j,gs}) - (\ln b^t_{j,gs} - \ln b^0_{j,gs})\right)}, \tag{23}$$

where $(\ln b^t_{j,gs} - \ln b^0_{j,gs})$ is given by Equation (22). The price index calculated based on these modified weights precisely reproduces the price index of equation (8) in Redding and Weinstein (2020), as shown in Appendix B.

Using sectoral import prices $\hat{P}^{M,t}_s$, that are given by Equations (19) and (20), we apply the OPI-M version of Equation (18) to calculate sectoral prices, $\hat{P}^t_s$. In our calculations, we utilize two different sets of weights:

1. FBW weights $\hat{d}^t_{j,gs}$ from equation (21).

2. RW weights $\hat{d}^t_{j,gs}$ from equation (23).

We refer to the price indices that use FBW weights as "OPI-M1", and those that use RW weights as "OPI-M2."

The resulting OPI-M indices have three components: the intensive margin, or the price changes within the common set; the extensive margin, or variety entry and exit; and the import-share adjustment, $(\hat{\pi}^t_{M,s})^{\frac{1}{\sigma_s - 1}}$, which corrects for changes in the relative price of imports.[8]

Aggregation. We now use sectoral prices $\hat{P}^t_s$ of equation (18), either OPI-D, or OPI-M1 and -M2 to construct price indices of food, $F = 1$, and non-food N. The latter is an aggregation of tradable sectors $s = \{2, ..., S-1\}$ and non-tradable services, $V = S$:

$$\begin{cases} \hat{P}^t_F = \hat{P}^t_1 & \text{[Food]} \\ \hat{P}^t_N = \left[\prod_{s=2}^{S-1} (\hat{P}^t_s)^{\beta_{s,t}}\right]^{1-\beta_{S,t}} \times (\hat{P}^t_S)^{\beta_{S,t}} & \text{[Non-food]} \end{cases} \tag{24}$$

Here, $\beta_{s,t}$ represents the share of consumer expenditure on sector s within the non-food tradable category, while $\beta_{S,t}$ is the expenditure share on non-tradable services. Both shares are allowed to vary over time t.

5 Data and Preliminary Analyses

In this section, we first briefly outline our data and parameter values, relegating the full details to the Data Appendix. We then show the salient features of our data, and conduct preliminary data analyses.

[8] In the literature, Feenstra (1994) and Broda and Weinstein (2006) focus on import price indices, and do not include import-share adjustments. Redding and Weinstein (2018), however, include this margin in their decomposition of the US aggregate price index.

5.1 Data Sources and Parameter Values: Outline

Micro Data for the Food Engel Curve. We use data from the Panel Study of Income Dynamics (PSID) to obtain household-level information on food consumption share and income. The PSID provides annual data for 1995-1996 and biannual data for 1997-2015. In addition, the PSID includes a rich set of household characteristics, such as age, education, and number of children, which we use as additional controls in Equation (9). All of these variables are available in the public version of the PSID. However, identifying regression (9) also requires region fixed effects. For this, we use confidential PSID data that provide the county geo-codes where households are located.

The remaining variables needed to estimate equation (9) are the prices of food and non-food, which vary by region and year. To capture regional price variation, we use data from the Cost of Living Index (COLI), compiled by the Council for Community and Economic Research (Council for Community and Economic Research, 2025). The COLI is based on prices of goods sampled at the metropolitan statistical area (MSA) level across the United States, which we map to the level of counties to augment with PSID information. We aggregate these prices to obtain COLI indices of food and non-food, expressed as regional deviations from the national average, which we denote by $P^t_{r,g,COLI}$ with g referring to food or non-food in region r and year t. We then combine these regional indices with national indices, either the OPI or CPI, to calculate $P^t_{r,F}$ and $P^t_{r,N}$. Specifically,

$$P^t_{r,g} = P^t_{r,g,COLI} \times P^t_g, \qquad g \in F, N, \tag{25}$$

where $P^t_{r,g,COLI}$ is the regional COLI index and P^t_g is the corresponding national OPI or CPI index. Our approach assumes that, within each year, regional differences in consumer prices are consistent across the various measured price indices.

Finally, we use the confidential PSID county codes to match households to the appropriate MSA-level measures from the COLI. This allows us to construct all the variables needed for estimating equation (9).

Macro Data for OPI-D and OPI-M We obtain OECD STAN data on gross production, exports and imports by sector by year. The STAN sectors are at the aggregate level of two-digit ISIC codes (International Standard Industrial Classification, version 4). Table 1 lists our 12 tradable sectors, one for the category of food, F, which aggregates over Agriculture and Manufacturing of Food, and 11 for non-food merchandise, G.

Our STAN data provide the values of the imported expenditure share, $\pi_s^{M,t}$ (*e.g.* equation 18), as imports divided by apparent consumption (*i.e*, gross output minus exports plus imports). The domestic expenditure share, $\pi_s^{D,t}$ is simply $1 - \pi_s^{M,t}$.

The food Engel curve is based on household consumption data, while products of some the sectors listed in Table 1 are not directly purchased by consumers (*e.g.* Mining). To address this concern, we obtain the consumption shares from the World Input-Output Database (WIOD) for the 11 non-food merchandise sectors in the category G. We then compute the relative consumption shares from WIOD, and multiply them by the CPI weight of category G to obtain β_s in OPIs

Table 1: Tradeable Sectors

	ISIC Code	Description
1	01-03 & 10-12	Food & Agriculture
2	13-15	Textile and Apparel
3	16-18	Wood and Paper
4	19	Refined Petroleum
5	20-21	Chemicals
6	22	Plastics
7	23	Minerals
8	24-25	Metals
9	26-28	Machinery and Electronics
10	29-30	Transport Equipment
11	31-32	Furniture & Other Mfg.
12	05-08	Mining

Notes: This table shows our food and non-food tradable sectors and their corresponding ISIC rev. 4 codes.

in Equation (24). For the aggregation into the non-food group of N, we obtain the category-level U.S. CPI weights from the BLS (Bureau of Labor Statistics), $\beta^t_{V,CPI}$ and $\beta^t_{G,CPI}$, and set $\beta_{S,t} = \beta^t_{V,CPI}/(\beta^t_{V,CPI} + \beta^t_{G,CPI})$.

Domestic Price, for OPI-D A challenge we face in bringing our formulas to data is that it is rare for researchers to directly observe the prices of domestic varieties,[9] from which the measure of domestic price inflation, $\hat{P}^{D,t}$, can be constructed. As a result, we use the price indices constructed and published by national statistical agencies. While a candidate for measuring domestic price inflation might seem to be CPI itself, note that CPI captures movements in import prices as well, leading to double counting.

The way we address this issue is motivated by our focus on comparing OPIs with statistical measures of price indices, like CPI itself. In this regard, we want a measure of domestic price inflation that can be thought of as the domestic component of the official CPI inflation for the groups of food and non-food, $\hat{P}^{D,t}_{CPI,g}$, $g = F, N$. Specifically we use sector-level import price indices provided by FRED and their corresponding weights to compute the statistical food and non-food import price indices, $\hat{P}^t_{IMP,g}$,$g = F, N$. We then recover $\hat{P}^{D,t}_{CPI,g}$ by assuming that, at the sector level, the CPI reflects a combination of import and domestic price indices, weighted by their respective import and domestic expenditure shares. We obtain:

$$\hat{P}^t_{D,CPI,g} = \left[\left(\hat{P}^t_{CPI,g} \right) \Big/ \left(\hat{P}^t_{IMP,g} \right)^{\pi^t_{M,g}} \right]^{\frac{1}{1-\pi^t_{M,g}}}, \tag{26}$$

where g indexes food (F) and non-food (N), and $\pi^t_{M,g}$ denotes the imported expenditure share of g in year t. Equation (26) says that if official import price indices, $\hat{P}^t_{IMP,g}$, measure inflation due to imports, then removing $\hat{P}^t_{IMP,g}$ from its corresponding CPI value gives us a measure of domestic

[9] An exception is Auer et al. (2022), who identify source countries from the product labels for a subsample of the universe of products in the Swiss market.

inflation, $\hat{P}^t_{D,CPI,g}$, which we use as domestic price change for food and non-food.[10]

Trade Data and Parameter Values, for OPI-M We draw on standard publicly available data for variety-level international trade (BACI-CEPII) information on unit values and expenditure shares on imported products, which we use in the construction of OPI-M1 and -M2. We then draw on the literature for the values of substitution elasticities: the goods-level substitution elasticity, σ_g in equations (20) - (23), and σ_s from equation (18), which is the sectoral substitution elasticity between imported and domestic bundles in OPI-M and the trade elasticity in OPI-D. For σ_g, we take the estimates from Broda and Weinstein (2006) at the level of 3-digit SITC (Standard International Trade Classification) codes. We have slightly over 200 SITC goods; *e.g.* footwear, metal containers for storage or transport, and food-processing machines.[11] We set σ_s as the mean values of $\{\sigma_g\}$ across goods g in sector s, following Imbs and Mejean (2015). As reported in the appendix, the sectoral elasticities (σ_s) range from about 2 for sectors such as Plastics, Electronics, and Machinery, to as high as 26 for Mining. These values are consistent with the range of estimates found in the trade literature.[12] We also experiment with alternative values of these elasticities, and obtain similar results as shown in Appendix A.3.1. Lastly, note that we do not require the σ^M_s values in our computation, as shown in equation (19). This is because we assume the set of goods is fixed as, unlike varieties, there is no entry or exit of goods.

5.2 Main Data Features

Domestic Shares Patterns of sector-level domestic expenditure shares play an important role in the construction of the OPI indices. Figure 2 plots the U.S. domestic expenditure share by sector by year, where the sectors are in the tradable categories of food and non-food merchandise. A salient feature is that US domestic expenditure shares over the period of 1995-2015 tended to decrease or stay unchanged, and they rarely increased. The decrease in domestic expenditure share has been particularly notable in Textile and Apparel, Chemicals, Plastics, Machinery and Electronics, and Furniture and Other Manufacturing. These decreases are likely driven by the global wave of trade liberalizations in the 1990's and early 2000's, such as China's rise starting in the early 1990s (*e.g.* Autor et al. 2013) and its subsequent WTO accession in 2001, the signing of NAFTA in 1994, and the expiration of the Agreement on Textiles and Clothing quotas in 2005, among other considerations.[13]

[10] Note that $\hat{P}_{IMP,g}$ is the import price index by the official U.S. statistical agency, and distinct from the OPI-M indices that we construct in this study.

[11] One may be concerned about the large number of parameters involved. In the Data Appendix, we show that we obtain similar results for OPI-M if we apply σ_s to all the goods within sector s. Note that the OPI-D indices do not require the parameters σ_g.

[12] As we use changes in disaggregated unit values in our analysis, our results are sensitive to outliers. We follow the previous literature in cleaning outliers in a theoretically-consistent manner without excluding them from our analysis (Redding and Weinstein, 2020). Details can be found in Data Appendix A.

[13] Additional examples of trade liberalizations involving the U.S. are as follows. The tariff cuts under the U.S.-Canada Free Trade Agreement were not complete until the end of 1998 (*e.g.* Lileeva and Trefler 2010), and the U.S. substantially lowered tariffs on Vietnamese products after the U.S.-Vietnam Bilateral Trade Agreement in 2001 (*e.g.* McCaig 2011).

Figure 2: Domestic Share by Sector, U.S.

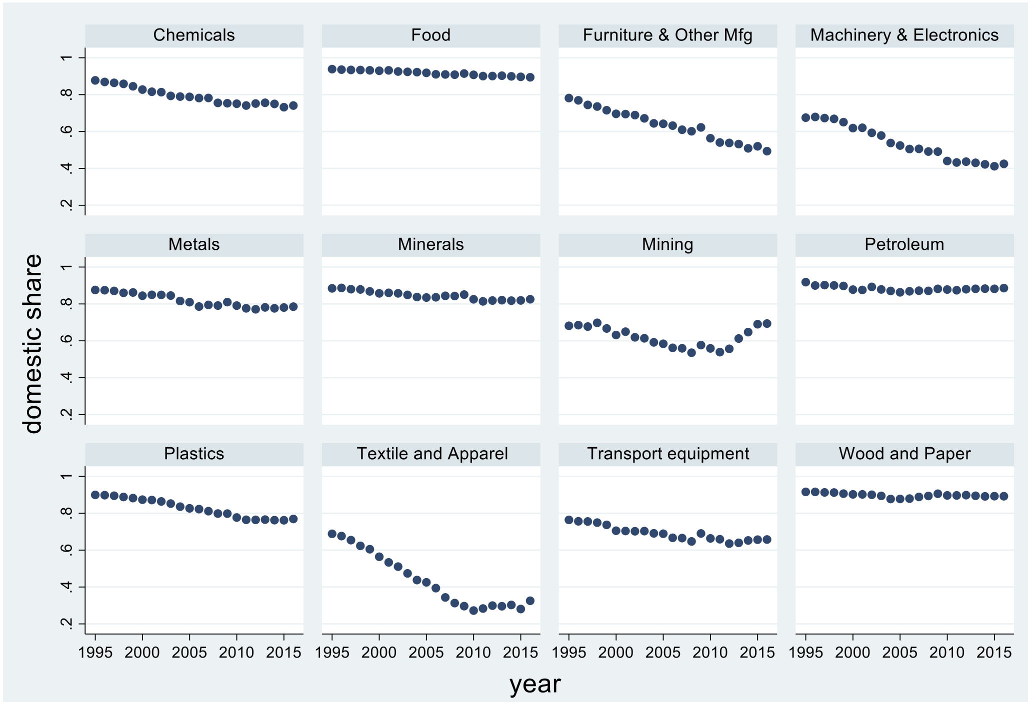

Notes: This figure shows the US domestic expenditure share for each industry between 1995 and 2015. The domestic share is the ratio of gross output minus exports to total absorbtion (which is gross output minus exports plus imports), the data of which come from STAN.

Empirical Engel Curves Our framework does not treat food and the other merchandise sectors symmetrically. Food is its own category and its own group (*e.g.* Figure 1), while other merchandise sectors are rolled together into a single category Non-Food. We choose to separate our sectors this way because of the well-documented nearly linear empirical relationship between food expenditures and log household income. This relationship is known as Engel's law,[14] which we illustrate in Figure 3. We divide households into deciles of residualized expenditure, and plot the mean residualized share of expenditures on food against the residualized mean log income. We plot residuals rather than levels as we would like to control for household characteristics such as the number of children and marital status. We see that food share decreases with log income, and that the relationship between food share and log income is close to a linear relationship. Figure 3 confirms that the Engel's law is a salient feature of our data.

6 Results

In this section, we report our results. We start with the preliminary results, in which we validate our approach of using the food Engel curve to estimate the bias of measured price indices. We then

[14]Moreover, the empirical relationship between log household income and the shares of sub-categories of food, or that between log income and shares of non-food products such as clothing, is not as stable across countries and time(*e.g.* Banks et al., 1997, Atkin et al., 2020).

Figure 3: Food Expenditure Share against Log Income

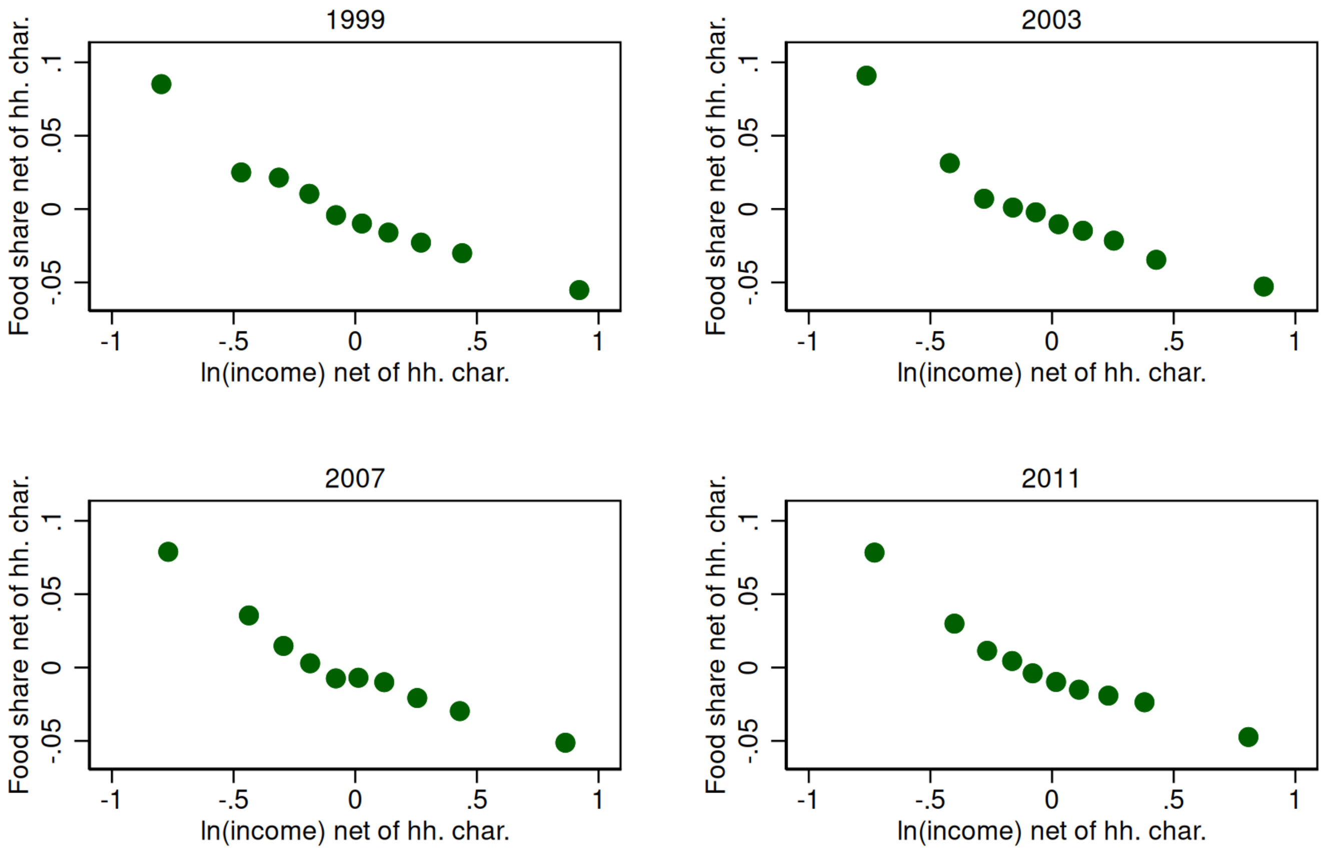

Notes: This figure shows the food expenditure share against log of expenditures along ten bins of the log expenditure distribution. Both food expenditure share and expenditures are residualized on household characteristics age, hours worked, education of the household head and spouse, number of children, and census region.

discuss our main results for OPI-M and OPI-D. Finally, we use OPI-M2, the OPI with the best fit of the food Engel curve, to show additional applications of our OPI indices in terms of welfare and simple counterfactuals.

6.1 Preliminary Results

As a first validation for our bias estimation method, we examine the estimates of the bias of the official U.S. CPI that our framework produces. Ever since the Boskin report (Boskin et al., 1997), there have been extensive studies of the U.S. CPI, and this literature has reached the consensus that the U.S. CPI is upward-biased. That is, using CPI as a deflator understates real income. Here, we are interested in comparing the bias estimates of the U.S. CPI from our analysis that is based on inference from the food Engel curve, with the estimates from this literature that are based on different approaches.

We use the official U.S. CPI to measure the prices of food and non-food in regression (9), and include the following household characteristics as controls: the number of children, the age, hours and education of the household head, and the age, hours and education of the spouse. We report the results in column (2) of Table 2. We see that the log-income coefficient, β_F in equation (9), is negative and significant. The statistical significance of the β_F estimates and the high R^2 show that the food Engel curve is an empirical regularity in household survey data.[15] Meanwhile, the year-dummy estimates under CPI are negative, and they are mostly statistically significant. These results provide evidence that our methodology correctly identifies the upward bias of the U.S. CPI, qualitatively. In addition, using equation (11), we are able to compute that the root mean squared bias, or RMSB, for U.S. CPI in our sample is 0.070. In other words, our estimates indicate that the U.S. CPI produces an upward bias of 0.70 log points on average per year. In comparison, Shapiro and Wilcox (1996) show that there is an 80% probability that the CPI bias lies between 0.6% and 1.5% per year, Gordon (2006) reports an annual bias of 0.8%, and Berndt (2006) reports 0.73%-0.9%. The similarity between our results and the literature provides evidence that our methodology also correctly identifies the upward bias of the U.S. CPI quantitatively.

Table 2: Food Engel Curve: CPI, Constant Price, and OPIs

Model	Ln(Income)	Ln(Pf/Pn)	Ln(Pf/Pn)^2
CPI	-0.0826***	0.0565***	-0.0650***
	(0.00138)	(0.00690)	(0.01950)
Constant	-0.0825***	0.0553***	-0.0636***
	(0.00138)	(0.00706)	(0.0192)
OPIM-1	-0.0825***	0.0540***	-0.0498**
	(0.00138)	(0.00789)	(0.01550)
OPIM-2	-0.0825***	0.0578***	-0.0446**
	(0.00138)	(0.00751)	(0.01650)
OPI-D	-0.0825***	0.0610***	-0.0587**
	(0.00138)	(0.00638)	(0.01850)
OPIG-1	-0.0825***	0.0588***	-0.0595**
	(0.00138)	(0.00664)	(0.01850)
OPIG-2	-0.0825***	0.0583***	-0.0602**
	(0.00138)	(0.00668)	(0.01860)

Notes: Each regression includes control variables for age, hours, and education of both head and spouse, and number of children, along with year dummies which are reported in Appendix Table A.3. Robust standard errors in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

For our second validation analysis, we use the constant price index in all years in the food Engel equation (9); $i.e$ we assume, naively, that $P_F^t = P_N^t = 1$ for all t. Because U.S. inflation is positive in most years, we expect that by using the constant price index as the deflator, we will overstate household real income. This implies that our year-dummy estimates under constant price index should be mostly positive.

To implement this analysis, we use household nominal income in the estimation of equation (9), and include the full set of control variables. Column (1) of Table 2 reports the results. The

[15]Previous studies listed in Footnote 1 that have estimated the food Engel curve for different time periods and different countries report their β_F estimates to be between -0.2 and -0.05. In turn, we will discuss the coefficient estimates of $\ln(P_F/P_N)$ and $\ln(P_F/P_N)^2$ in sub-section 6.2 below.

Figure 4: Root Mean Squared Bias by Price Index

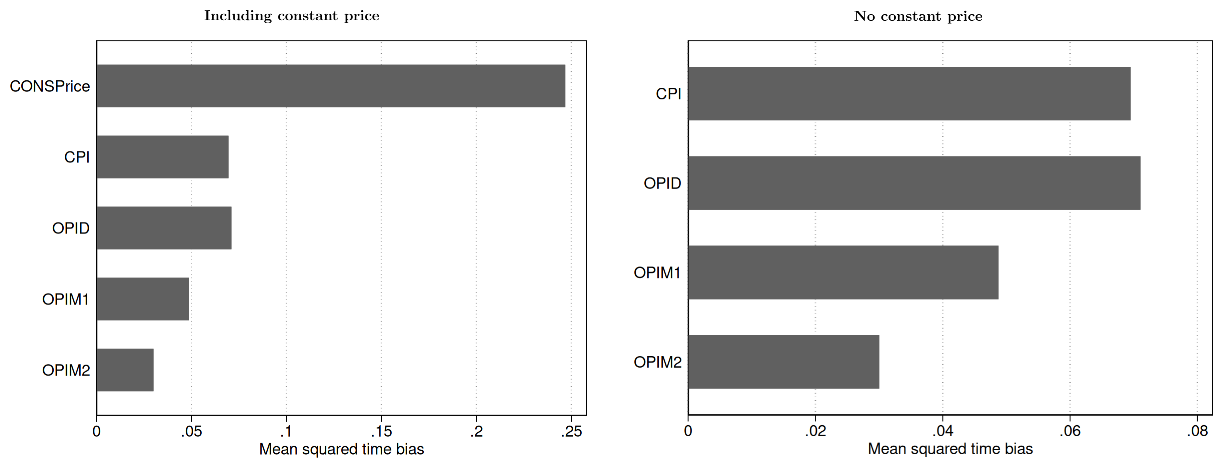

Notes: This figure presents estimates of root mean squared bias under different income deflators. The left panel shows results including constant price (treating nominal as real income), while the right panel omits constant price for ease of comparing the other price indexes.

estimated log-income coefficient, β_F, is similar to Column (2). The coefficient estimates of the year dummies are all positive in sign, and they are statistically significant starting 2003. These results provide evidence that our methodology correctly identifies the downward bias of the naive price index of constant prices. In addition, as illustrated in the left panel of Figure 4, the RMSB under constant prices is 0.247, over three times as high as the RMSB of CPI. This shows that our methodology also correctly ranks CPI as a more preferable price index than constant prices.

6.2 Results for OPI-M

Figure 5: Price indexes 1995-2015

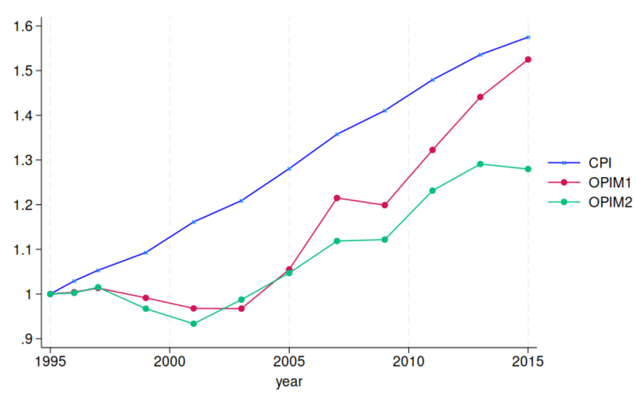

Notes: This figure shows all price indexes used in our analysis. Prices are normalised to one in 1995. We only include years in which we have PSID data on household expenditures. Before 1997, PSID data is available every year, but after 1997, we have only odd years.

Overall Price Level. We begin by displaying the aggregated OPI-M measures along with CPI in Figure 5. In this figure, "OPI-M1" is the weighted geometric mean of the food and non-food prices given by equation (24), where the weights are CPI weights, and likewise for "OPIM-2". While the aggregation is not used in our formal analyses with the food Engel curve estimation, it provides a concise way to show the differences between OPIs vs. CPI and to highlight the main features of our analytical framework.

The OPI-Ms incorporate gains from trade into over-time price changes by taking advantage of rich product-level trade data, and we see, from Figure 5, that both aggregated OPI-M measures follow different paths vs. CPI. First, both aggregated OPI-M measures are below CPI. This is consistent with the core message of the gains-from-trade literature that consumers experience lower price increases once gains from trade are taken into account. Second, the aggregated OPI-M2 measure is below the aggregated OPI-M1 measure, and this is consistent with Redding and Weinstein (2020). Finally, we also see that, strikingly, prices fell, relative to 1995, throughout the late 1990's and the early 2000's, according to both aggregated OPI-M measures. This likely reflects the expansion of global trade that we previously discussed for Figure 2.

Figure 5 clearly shows that the aggregated OPI-M measures are different from CPI, but does not clarify whether such differences are "good" or "bad". Specifically, while the lower levels of the aggregated OPI-M measures could help alleviate the CPI's upward bias, on average, they may also introduce downward biases, as illustrated by the naive price index of constant prices that we discussed in the previous sub-section. What empirical evidence do we need, to conclude that the OPI-M indices may be preferable to CPI, overall? What other empirical evidence do we need, to respect the striking claim that, during the late 1990's and early 2000's, prices in the U.S. may have actually decreased, because of gains from trade? To help answer these questions, we turn to our analyses using the food Engel curve.

Estimated Parameter Values Rows (3) and (4) of Table 2 report the results of the estimation of regression (9) with OPI-M1 and OPI-M2. In our data, the variation across households in income, demographics and food consumption identifies the coefficient of log income. Our data also has variation across county-by-year in the prices of food and non-food, and this variation is rich enough to separately identify the coefficients of log relative food price and its square, an innovation relative to previous studies. All coefficient estimates are statistically significant. Parameter estimates do not change much across specifications in Table 2, because the change of national-level indices from CPI to OPI-M does not affect the cross-household variation in our data, nor does it affect the county-by-year variation much, as shown by equation (25).

Our estimation results in Table 2 allow us to recover all the key parameters of the AIDS preferences of (2). It is reassuring that our estimates of Engel curve coefficients are similar across specifications. We get similar preference parameters regardless of which set of estimates we employ. For example, for OPI-M2, the estimated coefficient of log income is $\hat{\beta}_F = -0.0825$, and using the coefficients of relative price and its square, equation (12) implies that $\hat{\gamma} = -1.0812$, $\hat{\alpha}_F = 13.8062$.

We now discuss the intuition of these parameter values. First, we have discussed our β_F estimate in the previous sub-section. Second, β_F, γ and α_F all affect household decisions on the margin. Using equations (5) and (6), we obtain that the average household in our data has the own price elasticity of -0.50 and the income elasticity of 0.33 (see Appendix B for the details). These elasticity estimates are reasonable and comparable to the literature.[16] Finally, while only β_F enters into the calculation of the bias in equation (10), γ and α_F are jointly estimated with β_F in regression (9), and so β_F, γ and α_F all matter for the bias. In Section 6.4, we use these estimates to show how they can be used for counterfactual analyses.[17]

Figure 6: Year Dummies Over Time Across Price Indexes

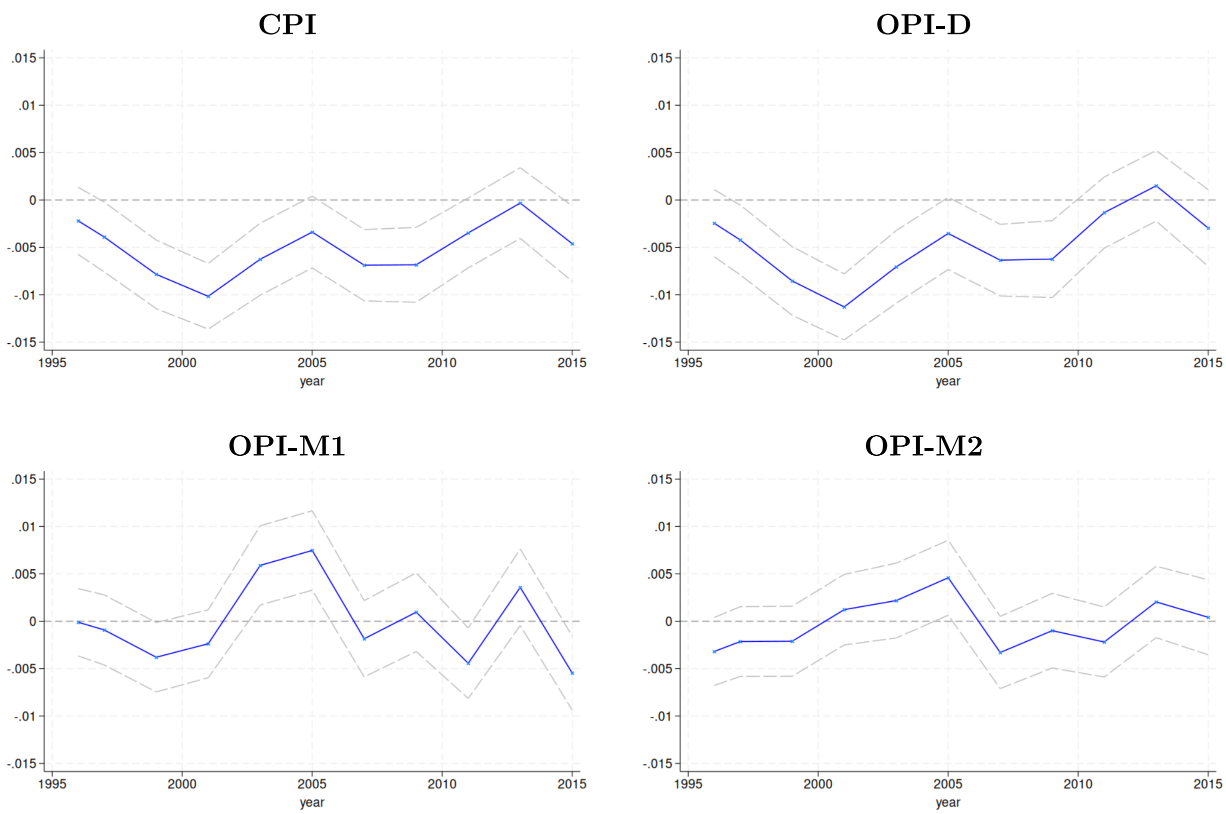

Notes: This figure presents the year dummies in Appendix Table A.3. The data period is 1995 to 2015. The year dummies are normalized to zero in 1995. We include 95% confidence intervals.

Ranking Price Indices We plot the estimated year dummies from Table 2 by price index for each year in Figure 6, as well as their 95 percent confidence intervals. We already reported the RMSB of these price indices in Figure 4, where in the right panel, we remove constant prices, to

[16] *e.g.* Nakamura et al. (2016) estimate that the income elasticity ranges from 0.6-0.7 and the own price elasticity ranges from -0.8 through -0.6.

[17] Note that the identification of α_0 does not matter for identifying the biases in measured price indices, which is what we are examining here. This is because the level of α_0 does not affect household decisions on the margin. This feature, however, makes it challenging to identify α_0 from a regression like (9), as noted in the literature, *e.g.* Almås et al. (2018). However, α_0 drives the level of model predicted household food consumption share. Because we identify all the other AIDS parameters, we can choose α_0 by matching the level of food expenditure shares.

highlight the comparison between OPI-M1 and OPI-M2 vs. CPI. The main features of Figures 4 and 6 are as follows.

First, as described above, the year dummies with CPI are consistently negative and they are significantly different from 0 for 7 out of 11 years, indicating that it overstates prices and understates real income. Second, with OPI-M1, the year dummies tend to be smaller in magnitude than with CPI, and they are statistically different from 0 for 5 out of 11 years. The RMSB shrinks by 30% relative to CPI, from 0.070 to 0.049. Lastly, with OPI-M2, the year dummies are statistically indistinguishable from 0 for 10 out of 11 years, and the RMSB shrinks by an additional 39% relative to OPI-M1, to 0.030.

These results show that the over-time price changes computed from product-level trade data, OPI-M, tend to track the empirical patterns of U.S. household consumption choices better than official CPI statistics. Because OPI-M incorporates gains from trade, our results provide evidence that U.S. households behave as if they take such gains into account when they make consumption decisions. In particular, the year dummies of OPI-M2 for 1999, 2001 and 2003 are all statistically indistinguishable from 0, implying that the U.S. households behaved in these years as if the gains from trade were so large that the overall price level decreased relative to 1995, as ilustrated in Figure 5.

The Components of OPI-M Our best performing price index is OPI-M2. OPI-M2 involves three main components: changes in unit values of existing varieties (intensive margin), changes in the set of varieties being imported (extensive margin), and a correction for unobserved domestic production (import share). We now discuss how these components contribute to its good performance.

First, Section 4 shows that OPI-M1 and -M2 have exactly the same extensive margin and the same import-share adjustment, but OPI-M2 incorporates the RW correction for the intensive margin that is designed to better measure taste changes and quality changes over time.[18]. Therefore, the RW correction is solely responsible for the better performance of OPI-M2 vs. OPI-M1, reducing the number of statistically significant year dummy estimates from 5 to 1, and then shrinking the RMSB by 39%. This finding is complementary to the result in Redding and Weinstein (2020) that the RW procedure helps correct the upward bias in the FBW procedure.

We then show the contribution of the extensive-margin component, by dropping it from OPI-M2, and then use this hypothetical index without variety entry and exit in regression (9). From Figure 7, we see that the removal of variety entry and exit from OPI-M2 more than doubles the RMSB, from 0.030 to 0.072, which is higher than the RMSB of CPI.[19] These findings are consistent with the previous studies that emphasize new foreign varieties as an important contributor to gains from trade (*e.g.* Feenstra (1994); Broda and Weinstein (2006)).

[18] Redding and Weinstein (2020) measure varieties as bar codes and interpret their results to reflect relative taste changes. Our results may reflect other relative changes, such as quality, because our variety measure of HS6 by exporter is broader than bar codes.

[19] We show the biases implied by the year dummy estimates in Appendix Figure A.2, to keep our exposition concise.

Next, we examine the contribution of the import-share-adjustment component, by, again, dropping it from OPI-M2. The resulting index is for import prices only, and we use it in the estimation of regression (9). The year dummy estimates are positive and mostly statistically significant, suggesting that the import price component of OPI-M2 understates prices and overstates real income (see Appendix Figure A.2). In addition, Figure 7 shows that the downward biases are so large that the RMSB almost triples, to 0.081, which, again, is larger than the RMSB of CPI. These results suggest that the import price component of OPI-M2 is overly optimistic about gains from trade. To see the intuition, recall that domestic expenditure shares for the US sample generally decrease over time (Figure 2), and so, imported expenditure shares generally increase. This implies that the relative price of imports decreases over time, and so the import price component of OPI-M2 is lower than OPI-M2. Our results thus provide empirical evidence that looking at the import price index alone may exaggerate gains from trade.

Finally, we further clarify the contribution of the intensive-margin component by removing it from OPI-M2, and then use this hypothetical index without any price data in regression (9). From Figure 7, we see that the RMSB increases by more than a factor of 5, to 0.169, which is more than twice as large as the RMSB of CPI. This finding shows that the product-level price data embedded in the intensive margin component is more important for OPI-M2 than the extensive-margin component and the import-share-adjustment component. While it is a simple point that a good price index needs good prices data, this point is important for interpreting our results for OPI-D and its extension, OPI-D2, below.

Figure 7: Mean Squared Bias by OPI-M2 Components

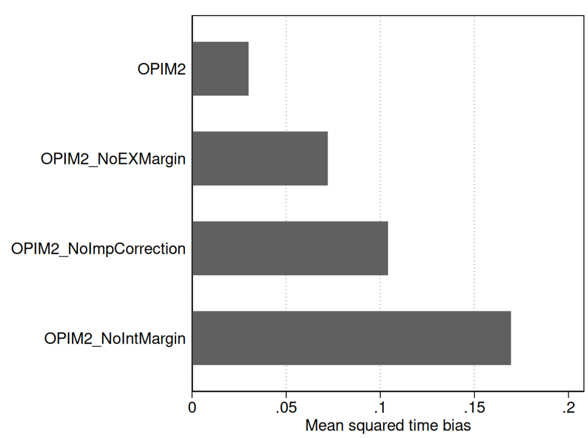

Notes: This figure presents the average bias in OPI-M2, leaving out each of the three OPI-M components one by one. These components are the correction for new and exiting varieties, changes in the import share of the economy, and observed change in prices for continuing varieties. The data period is 1995 to 2015.

Relating OPI-M to the CPI literature We have demonstrated, so far, that both OPI-M1 and OPI-M2 are more consistent with the patterns of food consumption by U.S. households than official CPI. We now clarify the intuition of these differences between CPI and OPI-M. We begin by briefly discussing how CPI is constructed in the United States. The U.S. constructs its CPI using a two-tier structure. The upper tier consists of about 270 ELI's (Entry Level Items), and the lower tier consists of individual items within ELI. The prices of individual items are collected by Bureau of Labor Statistics employees at retail outlets, and quantity data are available at the ELI-level, but not at the individual-item-level (*e.g.* Klenow and Kryvtsov, 2008). Therefore, item-level prices are aggregated into ELI-level via simple geometric mean, and then across ELI's via weighted geometric mean (*e.g.* Nakamura and Steinsson, 2008).

The resulting CPI index may deviate from the true prices that households use in their consumption decisions for the following reasons. First, an individual item may be replaced by a higher-quality item. This is not always easy to spot, and even if spotted, could be difficult to correctly adjust (*e.g.* Moulton, 1996). Second, new goods and new varieties may fail to show up in the CPI sample, and even when they do, the CPI procedure may not fully capture their effects on true household prices (*e.g.* Hausman, 2003). These two issues are widely recognized in the studies of U.S. CPI, commonly referred to as the quality bias and new-goods bias in that literature.

On the other hand, the OPI-M indices are explicitly designed to capture the effects of new varieties[20] and quality improvements if they originate in foreign countries. Consider one of the tradable sectors in the categories of food, F, or tradable merchandise, G. Products in this sector are differentiated into a number of varieties, some sourced from abroad.

Suppose, first, that some foreign country experiences a trade liberalization, and starts to export new varieties, and that the prices of existing varieties remain unchanged. In this scenario, consumers gain from trade, and the true price index they face decreases. While CPI may have difficulty with the new-variety effect, the extensive-margin component of OPI-M is explicitly designed to measure this effect. Suppose, alternatively, that some foreign country increases the quality of its existing varieties. In this scenario, consumers again gain from trade and experience lower prices, adjusted for quality. While CPI may have difficulty with the quality effect, this effect may be captured in the intensive-margin component of OPI-M2 via equation (22), because the import-value shares of the higher-quality existing varieties are likely to increase.[21] Together, the fact that OPI-M is able to capture both new goods and quality improvements originating from abroad – potentially a large source of U.S. consumer welfare gains in recent years – explains why OPI-M better tracks the empirical patterns of U.S. households' food consumption choices than CPI.

6.3 Results for OPI-D

We start with the aggregated OPI-D measure, whose construction is analogous to the aggregated OPI-M measures discussed in the previous subsection. Figure 5 shows that the aggregated OPI-D measure closely follows the path of the CPI. Note that OPI-D uses the domestic component of the CPI to measure domestic prices, as shown in equation (26). This component is higher than the CPI itself because the official U.S. import price index is lower than the CPI. However, OPI-D adjusts the

[20] Our terminology of "new variety" is in the context of trade models, and corresponds to both "new good" and "new variety" in the CPI literature.

[21] We continue to assume, as in section 4, that the over-time taste and quality changes within the common set, across all imports, are zero on geometric average.

domestic price using the domestic expenditure shares, and this adjustment pulls down the overall index because domestic shares tend to decrease over time due to expanding trade, as we showed earlier. In the end, the effect of using the domestic price component slightly dominates, and so the aggregated OPI-D measure is slightly higher than the CPI, as shown in Figure 5.

Column (5) of Table 2 reports the results of regression (9) with OPI-D. The estimated coefficients of log income, log relative food price and (to a slightly lesser extent) its square are all similar to previous price indices, in columns (2) through (4). The estimated year dummies and their 95 percent confidence intervals, on the other hand, are plotted in Figure 6. We see that, while these estimates are larger than those with CPI in magnitude, the differences are small, and they are statistically different from 0 for 6 out of 11 years. As shown in Figure 4, the RMSB of OPI-D is 0.072, slightly larger than the RMSD of CPI, and much larger than the RMSB of OPI-M. This is not driven by the adjustment using domestic expenditure shares, because these shares are symmetric to the import-shares-adjustment component of OPI-M, as shown in equation (18). Instead, it is driven by the use of the domestic price component of CPI. As we discussed previously, the U.S. CPI has difficulty with the new-variety effect and quality effect, but the import price component of OPI-M is better able to capture these effects. As a result, the import price component of OPI-M better reflects the way U.S. households make consumption decisions than the domestic price component of OPI-D, and this explains why OPI-D has a much larger RMSB than OPI-M. As we discussed previously, a good price index needs good price data.

6.4 Trade Shocks and Cost of Living Increases

Our estimation identifies the parameters of the AIDS preferences, so we can also use OPI for back-of-the-envelope calculations for cost of living changes in simple counterfactuals. We use our best performing price index, OPI-M2. The only parameter that remains to be pinned down is α_0 which cannot be identified from our food Engel curve estimation, and is not required for the preceding analysis of the biases in prices. However, we need it to pin down welfare levels in this subsection. We specifically obtain $\hat{\alpha}_0 = -156.4$, by matching the model predicted 99th percentile of household food consumption share to the data.[22]

Motivated by recent political discussions in the United States, we study how a 10% universal increase in the price of imports would impact U.S. households' cost of living, holding both domestic prices and household income unchanged. Our baseline is 2015, one year before Donald Trump was elected for the first time, and we use the import price changes and the 2015 domestic expenditure shares to solve for the change in the sectoral price index, $\hat{P}_s^t$, using equation (18). This implies that households optimally shift consumption between imports and domestic goods. Finally, we use the counterfactual changes in $\hat{P}_s^t$ to compute the changes in OPI-M2 and then plug them into the AIDS preferences of equations (1) and (2), to obtain the increase in cost of living by household income,

[22] We fit a smooth local polynomial to food expenditure share as a function of income for our 1995 data. Then we use the predicted food expenditure share of the 99th percentile of income as our target for the model. The reason we use a top percentile as a target is that our Engel curve is steep (in line with the literature). If we target lower percentiles we predict negative food expenditure shares for the highest incomes.

which is the percentage change in income required to maintain the households' pre-shock utility at post-shock prices.

Our results suggest that OPI-M2 would increase by 0.9% for Food and 0.8% for Non-food.[23] The differences in price response happen because domestic expenditure shares are different in Food and Non-food. For the median U.S. household, income would need to increase by around 0.8% to maintain pre-shock utility, indicating a modest increase in the cost of living due to higher import prices. In addition, the effects vary by household income because of non-homothetic utility. Specifically, poor households are hurt slightly more than rich households because they consume a larger share of food and food prices rise slightly more than non-food prices. These differences across households, however, are small in magnitude, as shown in Appendix Figure A.3's top panel.

We also simulate the scenarios in which the 10% increase in import prices happens only for Food, or only for Non-food. If only food import prices rise, the cost of living for households three standard deviations below mean income rises by just over 0.2%, while households three standard deviations above mean income are nearly unaffected by the price change as they spend so little on food. For the median household, the change in cost of living is also small, as food makes up a small fraction of household expenditure. On the other hand, if the rise in import prices is only on non-food, then rich households would experience a larger increase in cost of living, which is the opposite pattern as compared with the previous two scenarios. This happens because rich households consume a larger share of non-food. We plot these results in Appendix Figure A.3's middle and bottom panels.

7 Extension: OPI-G

In this section, we explore whether our analytical framework can shed light on which specific model specifications in the gains-from-trade studies that use general-equilibrium models deliver predictions that track observed household consumption decisions more closely. We do so by constructing open-economy indices that correspond to the following two commonly used specifications in this literature.

Theoretical Framework The first model specification follows the same setting as Sections 3 and 4, except that it endogenizes all prices by making the additional assumption that $\hat{P}^D = \hat{w}$ for every variety in the Home country, where w is the nominal wage. This specification corresponds to the multi-sector model in Costinot and Rodríguez-Clare (2014), and we call the resulting open-economy price index OPI-G1, where "G" refers to general-equilibrium models. Specifically, we have

$$\hat{P}_s^t = \hat{w}^t (\hat{\pi}_{D,s}^t)^{\frac{1}{\sigma_s - 1}}. \qquad \text{[OPI-G1]} \qquad (27)$$

Comparing equation (27) with equation (18), we see that the only difference between OPI-G1 vs. OPI-D is that we have replaced the domestic price in OPI-D, $\hat{P}_{D,s}^t$, which is observed from the data, with its general-equilibrium model prediction of $\hat{w}$.

The second model specification follows the multi-sector model with input-output linkages à la

[23] The overall category price increases are small as the United States has large domestic expenditure shares.

Caliendo and Parro (2015), which adds the following element into the specification of OPI-G1: sector $s = \{1, ..., S\}$ uses the outputs of all the sectors in the economy, including itself, as intermediate inputs. Specifically, assume that the production cost of sector $s = \{1, ..., S\}$ is a Cobb-Douglas aggregate of domestic factor services, with share γ_s, and the intermediate-input bundle, with share $1 - \gamma_s$. The latter is, in turn, a Cobb-Douglas aggregate of the price indices of all the sectors in the economy with shares α_{ks}, $k = \{1, ..., S\}$ and $\sum_{k=1}^{S} \alpha_{ks} = 1$. The input-output table of the economy is then the $S \times S$ matrix $\mathbf{A}$, whose element on row $k-$column s is $(1-\gamma_s)\alpha_{ks}$. The sector-level price index of each sector $s = \{1, ..., S\}$ can be then expressed as:

$$\hat{P}_s^t = \hat{w}^t \left[\prod_{k=1}^{S} \left(\hat{\pi}_{D,k}^t \right)^{\frac{\tilde{\alpha}_{ks}}{\sigma_k - 1}} \right]. \qquad \text{[OPI-G2]} \tag{28}$$

Here, $\tilde{\alpha}_{ks}$ represents the elements of the Leontief-inverse matrix of the economy; *i.e* it is the row $k-$column s element of $(\mathbf{I} - \mathbf{A}^{\mathbf{T}})^{-1}$, where $\mathbf{I}$ is the identity matrix and $\mathbf{A}^{\intercal}$ is the transpose of the input-output matrix $\mathbf{A}$. The parameters $\tilde{\alpha}_{ks}$ reflect the importance of sector k as an intermediate input for the production of sector s. To see the intuition of $\tilde{\alpha}_{ks}$, suppose that import prices fall in sector k, such that $\hat{\pi}_{D,k}^t < 1$. Without input-output linkages, this change lowers sector k's own price index, P_k^t, by $\left(\hat{\pi}_{D,k}^t\right)^{\frac{1}{\sigma_k - 1}}$, as given in equation (27), and has no impact on the other sectors' price indices. With input-output linkages, however, the other sectors $s \neq k$ experience price decreases of $\left(\hat{\pi}_{D,k}^t\right)^{\frac{\tilde{\alpha}_{ks}}{\sigma_k - 1}}$, where $\tilde{\alpha}_{ks} > 0$, as they all use sector k's outputs as inputs. In addition, the effect on sector k's own price index is amplified to $\left(\hat{\pi}_{D,k}^t\right)^{\frac{\tilde{\alpha}_{kk}}{\sigma_k - 1}}$, where $\tilde{\alpha}_{kk} > 1$ captures the extent of the amplification, because sector k uses all the other sectors' outputs as inputs. Comparing equation (28) with (27), we see that the only difference between OPI-G2 vs OPI-G1 is the addition of input-output linkages in (28).

Equations (27) and (28) give the sectoral prices of OPI-G1 and -G2, and these prices can then be aggregated into the group level of food, F, and non-food, N, using equation (24).

Data and Empirical Results In order to construct the OPI-G indices, we interpret the variable of $\hat{w}^t$ as over time changes in the labor cost of production for the U.S. economy, normalized by over-time changes in total-factor productivity (TFP).[24] We thus obtain labor cost data from STAN and TFP data from the Penn World Table (PWT) for the U.S. by year. Additionally, we obtain the data for the input-output linkages from WIOD.

When we use OPI-G1 and OPI-G2 in the estimation of regression (9), we again obtain similar coefficient estimates for log income, log relative food price and its square (see Appendix Table A.3). We find that OPI-G1 has the RMSB of 0.104, while OPI-G2 has the smaller RMSB of 0.085. These results suggest that the input-output linkage is a useful model element of trade theory. Specifically, while input-output linkages are absent from the underlying model of OPI-G1, they are a central feature of OPI-G2. As we discussed earlier, input-output linkages can amplify the direct effect of

[24] That is, $\hat{w}^t$ is proportional with changes the wage rate, but inversely proportional with productivity.

trade shocks in a sector, lowering the price there, because all sectors in the economy use the outputs of that sector as inputs. Therefore, the addition of input-output linkages is solely responsible for reducing the RMSB of OPIG2 by 18% relative to OPI-G1. This finding implies that the additional model element of input-output linkages delivers gains-from-trade predictions that better track the data patterns of U.S households' food consumption choices.

Our results also show that the RMSB of both OPI-G indices is larger than CPI, and they highlight the conundrum we face in interpreting the model variable of $\hat{w}^t$ in the data. On one hand, our interpretation, over time changes of production costs, is consistent with the underlying models, but it creates the challenge in the data that the computation of OPI-G does not use any price data. Specifically, the use of production costs vs. domestic prices is solely responsible for increasing the RMSB of OPI-G1 by 44%, almost one half, relative to OPI-D. This challenge is the main reason that the RMSB of OPI-G is larger than CPI; *i.e* it is difficult to come up with a good price index without using price data. On the other hand, an alternative interpretation is that, because $\hat{w}^t$ moves with the change of a nominal variable over time and routinely normalized to 1 in the literature, it can be interpreted as general inflation in the data, such as CPI. This interpretation, however, faces the challenge in theory that, because "general inflation" is outside of the underlying models, it is unclear whether the use of CPI, or other common measures of general inflation in the data, is genuinely consistent with the models. Additionally, any mismeasurement in productivity growth influences $\hat{w}^t$, creating another empirical challenge. We hope that future work can address both challenges in the empirical implementation of OPI-G.

8 Conclusion

Recent years have witnessed rapid growth in the studies that predict how much consumers gain from international trade. In contrast, our knowledge remains limited as to what extent the gains-from-trade predictions by this literature are consistent with observed household consumption decisions. In addition, this literature has used a number of model specifications and generated a range of gains-from-trade predictions. It is unclear whether specific model elements render the overall predictions more consistent, or less consistent, with household consumption decisions.

In this paper, we incorporate the core predictions from this literature about real consumption into over-time changes in open-economy price indices, or OPI indices. We then use the food Engel curve, a strong empirical regularity in micro household survey data, to quantify the deviations of OPI indices from the true price index that households use to deflate their income, and compare them with the deviations of CPI. We show that the OPI indices tend to follow the true price index more closely than official CPI. We also show that the overall gains-from-trade predictions tend to better track household consumption decisions if we allow for multiple industries and input-output linkages in quantitative general-equilibrium models, or if we correct for demand residuals in the computation of import-price indices as in Redding and Weinstein (2020). Our results provide a validation check of the core predictions, as well as specific model elements and specifications, of the gains-from-trade literature.

More broadly, official CPI is extensively used for policy. While there is a broad agreement that CPI does not track household prices well over time due to quality and new-goods biases, there is no consensus about better alternatives. For example, Meyer and Sullivan (2009) subtract an ad-hoc 0.8% from the growth in U.S. CPI each year in their analyses of poverty[25] because they do not have a better alternative. In turn, our results establish a connection between the gains-from-trade literature and the research on official CPI. The sufficient-static approach (*e.g.*, ACR), an important part of the gains-from-trade literature, provides a solution to the quality bias and new-goods bias of official CPI, both of which are difficult problems for the CPI literature. By operationalizing the complementarity between sufficient statistics from commonly-used trade theories and official CPI in a transparent and tractable way, our OPI-D indices provide one potential alternative to be used alongside CPI for policy analysis and discussion.

References

Adão, R., Costinot, A., and Donaldson, D. (2025). Putting quantitative models to the test: An application to the us-china trade war. *The Quarterly Journal of Economics*, 140(2):1471–1524. 1

Almås, I. (2012). International income inequality: Measuring ppp bias by estimating engel curves for food. *American Economic Review*, 102(2):1093–1117. 1, 1

Almås, I., Beatty, T. K., and Crossley, T. F. (2018). Lost in translation: What do engel curves tell us about the cost of living? 5, 17

Amiti, M., Dai, M., Feenstra, R. C., and Romalis, J. (2020). How did china's wto entry affect us prices? *Journal of International Economics*, 126:103339. 1

Anderson, J. E. and Van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American economic review*, 93(1):170–192. 1

Antràs, P. and De Gortari, A. (2020). On the geography of global value chains. *Econometrica*, 88(4):1553–1598. 3

Argente, D. and Lee, M. (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association*, 19(2):913–952. 5

Arkolakis, C., Costinot, A., and Rodriguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1):94–130. 1, 3.1

Arkolakis, C., Demidova, S., Klenow, P. J., and Rodríguez-Clare, A. (2008). Endogenous variety and the gains from trade. *American Economic Review*, 98(2):444–50. 3

[25] Meyer and Sullivan (2009) use CPI-U-RS, which is very similar to what we have used, CPI-U, after the late 1990s.

Arkolakis, C., Ramondo, N., Rodríguez-Clare, A., and Yeaple, S. (2018). Innovation and production in the global economy. *American Economic Review*, 108(8):2128–73. 3

Atkin, D., Faber, B., Fally, T., and Gonzalez-Navarro, M. (2020). Measuring welfare and inequality with incomplete price information. Technical report, National Bureau of Economic Research. 1, 4, 14

Auer, R., Burstein, A., Lein, S. M., and Vogel, J. (2022). Unequal expenditure switching: Evidence from switzerland. Technical report, National Bureau of Economic Research. 5, 9

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review*, 103(6):2121–2168. 5.2

Bai, L. and Stumpner, S. (2019). Estimating us consumer gains from chinese imports. *American Economic Review: Insights*, 1(2):209–24. 1

Balistreri, E. J., Hillberry, R. H., and Rutherford, T. F. (2011). Structural estimation and solution of international trade models with heterogeneous firms. *Journal of international Economics*, 83(2):95–108. 3

Banks, J., Blundell, R., and Lewbel, A. (1997). Quadratic engel curves and consumer demand. *Review of Economics and statistics*, 79(4):527–539. 1, 14

Baqaee, D. R., Burstein, A. T., and Koike-Mori, Y. (2024). Measuring welfare by matching households across time. *The Quarterly Journal of Economics*, 139(1):533–573. 1

Berndt, E. R. (2006). The boskin commission report after a decade: After-life or requiem? *International Productivity Monitor*, 12:61. 1, 2

Borusyak, K. and Jaravel, X. (2018). The distributional effects of trade: Theory and evidence from the united states. *Available at SSRN 3269579*. 5

Boskin, M. J., Dulberger, E. R., Gordon, R. J., Griliches, Z., and Jorgenson, D. W. (1997). The cpi commission: Findings and recommendations. *The American Economic Review*, 87(2):78–83. 2, 6.1

Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety. *The Quarterly journal of economics*, 121(2):541–585. 1, 4, 8, 5.1, 6.2, **??**

Caliendo, L., Dvorkin, M., Parro, F., et al. (2015). Trade and labor market dynamics. *Federal Reserve Bank of St. Louis Working Paper Series*, (2015-009). **??**

Caliendo, L. and Parro, F. (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies*, 82(1):1–44. 1, 7

Costa, D. L. (2001). Estimating real income in the united states from 1888 to 1994: Correcting cpi bias using engel curves. *Journal of political economy*, 109(6):1288–1310. 1, 1, 2, 2.2

Costinot, A. and Rodríguez-Clare, A. (2014). Trade theory with numbers: Quantifying the consequences of globalization. In *Handbook of international economics*, volume 4, pages 197–261. Elsevier. 1, 7

Council for Community and Economic Research (2025). Cost of living index (coli). `https://coli.org/`. Accessed: 2025-06-04. 5.1

Deaton, A. and Muellbauer, J. (1980). An almost ideal demand system. *The American economic review*, 70(3):312–326. 1, 2.1

Deaton, A. S. and Muellbauer, J. (1986). On measuring child costs: With applications to poor countries. *Journal of Political Economy*, 94(4):720–744. 1

Du, X. and Wang, Z. (2021). Multinational production, innovation relocation, and the consequences of globalization. *Innovation Relocation, and the Consequences of Globalization*. 3

Eaton, J., Eslava, M., Jinkins, D., Krizan, C. J., and Tybout, J. R. (2021). A search and learning model of export dynamics. Technical report, National Bureau of Economic Research. 3

Eaton, J., Jinkins, D., Tybout, J. R., and Xu, D. (2022). Two-sided search in international markets. Technical report, National Bureau of Economic Research. 3

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779. 1

Engel, E. (1895). *Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt*. Heinrich. 1

Fajgelbaum, P. D. and Khandelwal, A. K. (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics*, 131(3):1113–1180. 1, 5

Fally, T. and Sayre, J. (2018). Commodity trade matters. Technical report, National Bureau of Economic Research. 3

Farrokhi, F. (2020). Global sourcing in oil markets. *Journal of International Economics*, 125:103323. 3

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, pages 157–177. 1, 1, 4, 8, 6.2

Feenstra, R. C. and Weinstein, D. E. (2017). Globalization, markups, and us welfare. *Journal of Political Economy*, 125(4):1040–1074. 3

General Accounting Office (1995). Us imports: Unit values vary widely for identically classified commodities. Technical report, Report GAO/GGD-95-90. A.1.1

Giri, R., Yi, K.-M., and Yilmazkuday, H. (2021). Gains from trade: Does sectoral heterogeneity matter? *Journal of International Economics*, 129:103429. 3, ??, A.3.2

Gordon, R. J. (2006). The boskin commission report: A retrospective one decade later. Technical report, National Bureau of Economic Research. 1, 2

Greenlees, J. S. and McClelland, R. (2011). New evidence on outlet substitution effects in consumer price index data. *Review of Economics and Statistics*, 93(2):632–646. 2

Hamilton, B. W. (2001). Using engel's law to estimate cpi bias. *American Economic Review*, 91(3):619–630. 1, 1, 2, 2.2

Hausman, J. (2003). Sources of bias and solutions to bias in the consumer price index. *Journal of Economic Perspectives*, 17(1):23–44. 1, 6.2

Hsieh, C.-T. and Ossa, R. (2016). A global view of productivity growth in china. *Journal of international Economics*, 102:209–224. 3

Imbs, J. and Mejean, I. (2015). Elasticity optimism. *American Economic Journal: Macroeconomics*, 7(3):43–83. 5.1, ??, A.3.2

Jaravel, X. and Lashkari, D. (2024). Measuring growth in consumer welfare with income-dependent preferences: Nonparametric methods and estimates for the united states. *The Quarterly Journal of Economics*, 139(1):477–532. 1

Jaravel, X. and Sager, E. (2025). What are the price effects of trade? evidence from the u.s. and implications for quantitative trade models. Working paper, March 2025. 1

Kehoe, T. J., Pujolas, P. S., and Rossbach, J. (2017). Quantitative trade models: Developments and challenges. *Annual Review of Economics*, 9:295–325. 1

Kehoe, T. J., Rossbach, J., and Ruhl, K. J. (2015). Using the new products margin to predict the industry-level impact of trade reform. *Journal of International Economics*, 96(2):289–297. 1

Klenow, P. J. and Kryvtsov, O. (2008). State-dependent or time-dependent pricing: Does it matter for recent us inflation? *The Quarterly Journal of Economics*, 123(3):863–904. 6.2, A.1.1

Krolikowski, P. M. and McCallum, A. H. (2021). Goods-market frictions and international trade. *Journal of International Economics*, 129:103411. 3

Levchenko, A. A. and Zhang, J. (2016). The evolution of comparative advantage: Measurement and welfare implications. *Journal of Monetary Economics*, 78:96–111. 3

Lileeva, A. and Trefler, D. (2010). Improved access to foreign markets raises plant-level productivity... for some plants. *The Quarterly journal of economics*, 125(3):1051–1099. 13

McCaig, B. (2011). Exporting out of poverty: Provincial poverty in vietnam and us market access. *Journal of International Economics*, 85(1):102–113. 13

Melser, D. and Syed, I. A. (2016). Life cycle price trends and product replacement: Implications for the measurement of inflation. *Review of Income and Wealth*, 62(3):509–533. 2

Meyer, B. D. and Sullivan, J. X. (2009). Five decades of consumption and income poverty. Technical report, National Bureau of Economic Research. 8, 25

Moulton, B. R. (1996). Bias in the consumer price index: what is the evidence? *Journal of Economic perspectives*, 10(4):159–177. 1, 6.2

Moulton, B. R., Moses, K. E., Gordon, R. J., and Bosworth, B. P. (1997). Addressing the quality change issue in the consumer price index. *Brookings Papers on Economic Activity*, 1997(1):305–366. 2

Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464. 6.2

Nakamura, E., Steinsson, J., and Liu, M. (2016). Are chinese growth and inflation too smooth? evidence from engel curves. *American Economic Journal: Macroeconomics*, 8(3):113–44. 1, 2.2

Ossa, R. (2015). Why trade matters after all. *Journal of International Economics*, 97(2):266–277. 3

Redding, S. and Weinstein, D. E. (2018). Accounting for trade patterns. *Princeton University, mimeograph*. 8

Redding, S. J. and Weinstein, D. E. (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for ces preferences. *The Quarterly Journal of Economics*, 135(1):503–560. 1, 4, 4, 4, 12, 6.2, 6.2, 18, 8, A.1.1, B.1

Shapiro, M. D. and Wilcox, D. W. (1996). Mismeasurement in the consumer price index: An evaluation. *NBER macroeconomics annual*, 11:93–142. 2

Simonovska, I. and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of international Economics*, 92(1):34–50. A.3.2

Vartia, Y. O. (1976). Ideal log-change index numbers. *scandinavian Journal of statistics*, pages 121–126. 1

Appendices for "Assessing the Aggregate Price Effects of Trade Using the Food Engel Curve"

Farid Farrokhi, David Jinkins, Chong Xiang

A Data Appendix

A.1 Additional Details about Data

A.1.1 BACI-CEPII data

Our disaggregated variety-level trade data come from BACI-CEPII, which has annual values and quantities of bilateral trade at the level of 6-digit Harmonized System (HS6) codes. We use BACI data to compute variety-level unit values and import-value shares, the variables $p^t_{j,gs}$ and $\lambda^t_{j,gs}$ in equations (20) - (23), by defining a variety j as a pair of export country and HS6 product.

While the computation of the OPI-M indices accommodates the change of the variety set within a given SITC good over time, it requires that the common set of varieties be non-empty. In addition, the computation of the intensive margin of OPI-M2 involves taking the simple geometric mean within the common set, which might be an issue if the common set contains too few varieties. For these cases, we thus merge certain SITC goods in the raw BACI data, to ensure that the common set contains at least 20 varieties for each SITC good by year. Table A.1 lists the SITC goods that are subject to this consideration in merging the data for the U.S.

The OPI-M indices incorporate gains from trade into over-time price changes by taking advantage of rich product-level trade data. In practice, the variety-level prices computed from such data are typically noisy (*e.g.* General Accounting Office (1995)), and our data is no exception. For example, our U.S. imports data cover 5,354 unique HS6 products exported by 218 countries, and over 89,000 varieties (export country by HS6) are present in both 1996 and the base year of 1995. Among these varieties, the absolute values of log price changes, over 1995-1996, have the mean of 0.76 and the standard deviation of 1.05. Following common practice in previous studies, *e.g.* Redding and Weinstein (2020), we exclude the varieties with large price changes from the computation of the average price changes within the common set in the intensive margin, and treat them as variety entry and exit in the extensive margin instead. Our idea is that, if the exclusion removes the varieties with noisy price changes, then the remaining varieties should show similar aggregate moments in their price changes as compared with high-quality price data. Using one such data set, the raw data used in the construction of U.S. CPI, Klenow and Kryvtsov (2008) show that the mean of 8-month changes in log prices in absolute value is 0.11. Through experimentation, we find that if the exclusion cutoff is set to the median changes in log prices in absolute value, relative to the base year, then the mean value of one-year log price changes (absolute values) in the common set in our data is 0.13, close to the value of 0.11.[26]

[26] We have experimented with other exclusion cutoffs and obtained larger mean values. *e.g.* it is 0.23 for the U.S. if we use the 75th percentile as the cutoff, and 0.34 if we use the 90th percentile.

Table A.1: Merging SITC Goods Featuring Too Few Varieties in the U.S.

Old SITC code	New SITC code	New SITC sigma	New SITC Name
23	22	4.71	MILK AND CREAM AND MILK PRODUCTS OTHER THAN BUTTER
41	47	2.2	CEREAL MEALS AND FLOURS, N.E.S.
43	47	2.2	CEREAL MEALS AND FLOURS, N.E.S.
231	292	1.55	CRUDE VEGETABLE MATERIALS, N.E.S.
244	246	1.92	WOOD IN CHIPS OR PARTICLES AND WOOD WASTE
247	246	1.92	WOOD IN CHIPS OR PARTICLES AND WOOD WASTE
261	265	1.98	VEGETABLE TEXTILE FIBERS (OTHER THAN COTTON AND JU
263	265	1.98	VEGETABLE TEXTILE FIBERS (OTHER THAN COTTON AND JU
264	265	1.98	VEGETABLE TEXTILE FIBERS (OTHER THAN COTTON AND JU
274	278	4.76	CRUDE MINERALS, N.E.S.
277	278	4.76	CRUDE MINERALS, N.E.S.
281	278	4.76	CRUDE MINERALS, N.E.S.
283	285	2.66	ALUMINUM ORES AND CONCENTRATES (INCLUDING ALUMINA)
284	285	2.66	ALUMINUM ORES AND CONCENTRATES (INCLUDING ALUMINA)
287	333	27.85	PETROLEUM OILS AND OILS FROM BITUMINOUS MINERALS,
289	285	2.66	ALUMINUM ORES AND CONCENTRATES (INCLUDING ALUMINA)
322	321	2.18	COAL, PULVERIZED OR NOT, BUT NOT AGGLOMERATED
325	321	2.18	COAL, PULVERIZED OR NOT, BUT NOT AGGLOMERATED
333	333	27.85	PETROLEUM OILS AND OILS FROM BITUMINOUS MINERALS,
342	335	2.79	RESIDUAL PETROLEUM PRODUCTS, N.E.S. AND RELATED MA
343	335	2.79	RESIDUAL PETROLEUM PRODUCTS, N.E.S. AND RELATED MA
344	335	2.79	RESIDUAL PETROLEUM PRODUCTS, N.E.S. AND RELATED MA
712	718	1.23	POWER GENERATING MACHINERY AND PARTS THEREOF, N.E.
751	751	2.8	OFFICE MACHINES
752	751	2.8	OFFICE MACHINES
761	751	2.8	OFFICE MACHINES
783	783	2.98	ROAD MOTOR VEHICLES, N.E.S.
785	783	2.98	ROAD MOTOR VEHICLES, N.E.S.

Notes: This table reports how we merge and aggregate SITC goods featuring a common set with less than 20 varieties in the U.S.

A.1.2 STAN data

We obtain OECD STAN data on gross production, exports and imports by sector by year. The STAN sectors are aggregate ISIC (International Standard Industrial Classification) version 4. We have 1 sector for the food category, F, which combines agricultural and manufactured food, and 11 sectors for the category of tradeable merchandise, G. Table 1 in the main body of the paper lists the brief descriptions and ISIC-revision-4 codes of our 12 tradeable sectors, one for the category of food, F, and 11 for non-food merchandise, G.

Our STAN data provide the values of sector-level domestic and imported expenditure shares, $\pi_s^{D,t}$ and $\pi_s^{M,t} = 1 - \pi_s^{D,t}$, in equation (18), with $\pi_s^{D,t}$ corresponding to gross output minus exports divided by apparent consumption, which is gross output minus exports plus imports.

We encounter missing gross output data for 1995 and 1996 in sectors 31-32 (furniture & other manufacturing) and 01-03 (food & agriculture) for the U.S. To address this issue, we regress gross output on values of imports and exports for each sector-by-country using the years with non-missing data. The R^2 of these regressions are high (*e.g.* it is 0.96 for U.S. 01-03. We extrapolate the missing gross output values using the regression coefficients and the data for imports and exports.

A.2 Final consumption and intermediate use

CPI and the food Engel curve are based on household consumption data, while many goods in the trade data, which are used to compute the OPI indices, are not directly purchased by consumers (*e.g.* iron ore, crude oil). To address this concern, we use the consumption shares from WIOD to compute sectoral weights, β_s, when we aggregate across the non-food merchandise sectors in the category G (*e.g.* equation 24).

We illustrate how WIOD based consumption shares differ from shares based on absorbtion or "apparent consumption" from STAN. To construct the apparent consumption measures, we obtain the sectoral shares within category G using total absorption (gross output minus exports plus imports), and then multiply these relative shares by the CPI weight of category G. While the WIOD- and STAN-based consumption shares differ for individual non-food merchandise sectors, they have the same cross-sector aggregate within G by construction.

Figure A.1: Sectoral Consumption Shares, WIOD Data vs. STAN Data, U.S.

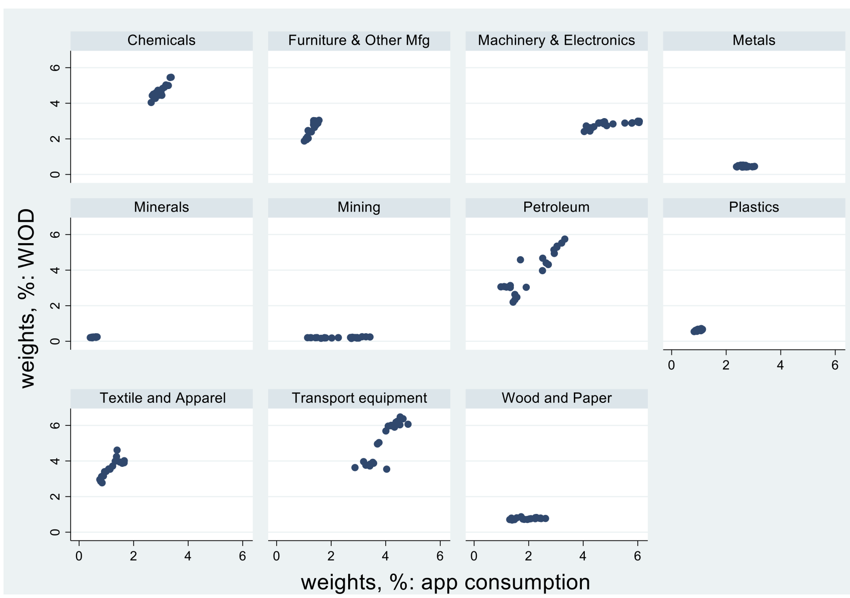

Notes: This figure shows the weights of each of the eleven non-food tradeable sectors.

Figure A.1 plots β_s, based on WIOD data, against the STAN-based consumption shares by year by sector, for the U.S. Mining has sizable apparent consumption and its STAN-based consumption shares range from roughly 2% to 4%. However, most of its goods are intermediate goods, and so its WIOD-based consumption shares are close to 0. We see a similar pattern for the sectors of Wood and Paper and Metals. This implies that the sectors for which intermediate goods are less important, such as Textile and Apparel and Transport Equipment, have larger WIOD-based consumption shares than STAN-based ones. Overall, Figure A.1 illustrates that WIOD data assign low consumption shares to the sectors largely producing intermediate goods.

A.2.1 Other Data

Our data on input-output linkages come from WIOD (World Input-Output Database). We use WIOD to compute the Leontief-inverse matrix, and obtain the values of its elements, the variables $\tilde{\alpha}_{ks}$ in equation (28).

A.3 Additional Details about Parameter Values

A.3.1 Industry-Mean Substitution Elasticity for OPI-M

As discussed in Sections 4 and 5, we use goods-specific substitution elasticities, σ_g, to compute the OPI-M indices, and one may be concerned about the large number of parameter values involved. In this sub-section, we construct the OPI-M indices with sector-level substitution elasticities, σ_s, instead. To do so, we maintain the same goods layers within sectors as before, and apply σ_s to the goods-level price indices. This change affects the extensive margins of both OPI-M indices, as shown in equation (20), as well as the intensive margin of OPI-M2, as shown in equation (20) and (22)-(23). Meanwhile, this change has no effect on the intensive margin of OPI-M1, as shown in equations (20) and (21), or the import-share adjustment of both indices, as shown in equations (18).

We obtain similar year-dummy estimates and similar average biases. The correlation coefficient for year-dummy estimates across OPI-M1 and -M2 is 0.78. The RMSB of OPI-M1 is 0.047 (vs. 0.049 in our main specification), and that of OPI-M2 is 0.035 (vs. 0.030).

A.3.2 Industry-level Trade Elasticities across the Literature

Table A.2 lists the values of the trade elasticity, $\theta_s = \sigma_s - 1$, as we use in this paper, and those used in a few studies in the literature. It also reports the mean values of θ_s for each study, as well as the correlation coefficients of these vectors of θ_s.

Table A.2: Trade Elasticity Estimates across Select Papers

Sectors	Ours, based on BW	Imbs and Mejean (Feenstra)	Giri et al (2021)	Imbs and Mejean (CP)	Caliendo and Parro (2015)
Food & Agriculture	2.9	6.1	3.6	4.0	2.6
Textile and Apparel	2.4	6.0	4.4	5.7	5.6
Wood and Paper	1.6	3.2	3.6	12.4	10.0
Refined Petroleum	8.0	8.5		41.8	51.1
Chemicals	2.1	4.2	3.8	5.2	4.8
Plastics	0.9	3.5	4.1	2.4	1.7
Minerals	1.0	4.7	5.1	1.7	2.8
Metals	4.5	3.8	7.0	13.1	6.1
Machinery and Electronics	1.0	6.0	3.3	12.5	8.2
Transport Equipment	2.3	4.9	4.5	5.5	0.7
Furniture & Other Mfg.	1.0	3.6	4.5	7.7	5.0
Mining	25.6	9.6		24.5	15.7
Mean	4.4	5.3	4.4	11.4	9.5
Correlation Matrix		0.8	0.6	0.6	0.4
			0.4	0.7	0.6
				0.4	0.2
					0.9

Notes: This table reports each industry's trade elasticity estimates, θ_s, as we use in this paper based on "BW" corresponding to (Broda and Weinstein, 2006), and across select papers in the literature: (Imbs and Mejean, 2015) based on Feenstra's procedure and based on CP, Caliendo and Parro's procedure, Giri et al. (2021), and (Caliendo et al., 2015). In case that two or multiple industries in another paper map to one industry in ours, we use the average across those multiple industries. The table, also, shows the mean value of estimates in each column, as well as the correlation matrix between any pair of columns.

The mean of our θ_s is 4.43, close to the often cited benchmark value of 4 in Simonovska and Waugh (2014). It is also similar to the mean values in Imbs and Mejean (2015) (Feenstra procedure) and Giri et al. (2021). In addition, our vector of θ_s is fairly well correlated with the ones in literature, with the correlation coefficients between our θ_s and the literature's θ_s being comparable to those among the literature's θ_s.

B Theory Appendix

B.1 Formulas for OPI-M

Nested CES Structure. Sector-level consumption bundle aggregates imported and domestic bundles in the following CES fashion:

$$C_s^t = \left[\left(C_s^{M,t}\right)^{(\sigma_s-1)/\sigma_s} + \left(C_s^{D,t}\right)^{(\sigma_s-1)/\sigma_s} \right]^{\sigma_s/(\sigma_s-1)}$$

where σ_s is the elasticity of substitution between imported and domestic sector–level bundles. In turn, the imported bundle of sector s, $C_s^{M,t}$, aggregates over imported bundle of goods within sector

s,

$$C_s^{M,t} = \left[\sum_{g \in \Omega_s} \left(C_{gs}^{M,t}\right)^{(\sigma_s^M - 1)/\sigma_s^M} \right]^{\sigma_s^M/(\sigma_s^M - 1)}$$

where σ_s^M is the elasticity of substitution between imported bundle of goods within sector s. The imported bundle of good g in sector s aggregates over varieties j,

$$C_{gs}^{M,t} = \left[\sum_{j \in \Omega_{gs}^{M,t}} \left(c_{j,gs}^t\right)^{(\sigma_g - 1)/\sigma_g} \right]^{\sigma_g/(\sigma_g - 1)}$$

where σ_g is the elasticity of substitution between varieties within good g.

Weights in Intensive Margin of the Price Index. Below, we derive the price index that corresponds to the consumption bundle of the lowest tier, $C_{gs}^{M,t}$. Derivations for upper tiers are similar. To simplify the notation, below we drop superscript M, and subscript gs, and work with:

$$C^t = \left[\sum_{j \in \Omega^t} \left(b_j^t\right)^{1/\sigma} \left(c_j^t\right)^{(\sigma - 1)/\sigma} \right]^{\sigma/(\sigma - 1)} \tag{B.1}$$

The associated price index is:

$$P^t \equiv P(\{p_j^t\}, \{b_j^t\}, \Omega^t) = \left[\sum_{j \in \Omega^t} b_j^t (p_j^t)^{1-\sigma} \right]^{1/(1-\sigma)} \tag{B.2}$$

Share of expenditure on variety j equals:

$$\lambda_j^t \equiv \lambda(\{p_j^t\}, \{b_j^t\}, \Omega^t) = \frac{b_j^t (p_j^t)^{1-\sigma}}{P_t^{1-\sigma}} \tag{B.3}$$

And, the change to the price index is defined as:

$$\hat{P}^t = \frac{P(\{p_j^t\}, \{b_j^t\}, \Omega^t)}{P(\{p_j^{t-1}\}, \{b_j^{t-1}\}, \Omega^{t-1})} \tag{B.4}$$

We only focus on the intensive margin of price changes since our derivations for the extensive margin is precisely the same as the ones in the previous literature. Focusing on the intensive margin, suppose the set of varieties purchased in t and $t-1$ are the same, $\Omega^t = \Omega^{t-1} = \Omega$. By equation (B.3), $P^t = (\lambda_j^t / b_j^t)^{1/(1-\sigma)} p_j^t$. Using this expression and the definition (B.4),

$$\hat{P}^t = \Big(\frac{\lambda_j^t / b_j^t}{\lambda_j^{t-1} / b_j^{t-1}} \Big)^{1/(\sigma - 1)} \Big(\frac{p_j^t}{p_j^{t-1}} \Big)$$

Taking logs:

$$\ln \hat{P}^t = \frac{1}{\sigma - 1} \ln\left(\frac{\lambda_j^t / b_j^t}{\lambda_j^{t-1}/b_j^{t-1}}\right) + \ln\left(\frac{p_j^t}{p_j^{t-1}}\right)$$

By reorganizing,

$$\frac{1}{\ln\left(\frac{\lambda_j^t/b_j^t}{\lambda_j^{t-1}/b_j^{t-1}}\right)} \ln \hat{P}^t - \frac{1}{\ln\left(\frac{\lambda_j^t/b_j^t}{\lambda_j^{t-1}/b_j^{t-1}}\right)} \ln\left(\frac{p_j^t}{p_j^{t-1}}\right) = \frac{1}{\sigma - 1}$$

Multiply this equation by $(\lambda_j^t - \lambda_j^{t-1})$ and sum over varieties $j \in \Omega$,

$$\sum_{j\in\Omega} \frac{\lambda_j^t - \lambda_j^{t-1}}{\ln\left(\frac{\lambda_j^t/b_j^t}{\lambda_j^{t-1}/b_j^{t-1}}\right)} \ln \hat{P}^t - \sum_{j\in\Omega} \frac{\lambda_j^t - \lambda_j^{t-1}}{\ln\left(\frac{\lambda_j^t/b_j^t}{\lambda_j^{t-1}/b_j^{t-1}}\right)} \ln\left(\frac{p_j^t}{p_j^{t-1}}\right) = \sum_{j\in\Omega} (\lambda_j^t - \lambda_j^{t-1}) \frac{1}{\sigma - 1}$$

From $\sum_{j\in\Omega} \lambda_j^t = \sum_{j\in\Omega} \lambda_j^{t-1} = 1$, it follows that $\sum_{j\in\Omega}(\lambda_j^t - \lambda_j^{t-1}) = 0$. Therefore,

$$\ln \hat{P}^t \sum_{j\in\Omega} \frac{\lambda_j^t - \lambda_j^{t-1}}{\ln(\lambda_j^t/b_j^t) - \ln(\lambda_j^{t-1}/b_j^{t-1})} = \sum_{j\in\Omega} \frac{\lambda_j^t - \lambda_j^{t-1}}{\ln(\lambda_j^t/b_j^t) - \ln(\lambda_j^{t-1}/b_j^{t-1})} \ln\left(\frac{p_j^t}{p_j^{t-1}}\right)$$

Hence, the change to price index equals:

$$\ln \hat{P}^t = \sum_{j\in\Omega} d_j^t \ln\left(\frac{p_j^t}{p_j^{t-1}}\right)$$

where the weight on variety j, d_j^t , is given by:

$$d_j^t = \frac{\frac{\lambda_j^t - \lambda_j^{t-1}}{\ln(\lambda_j^t/b_j^t) - \ln(\lambda_j^{t-1}/b_j^{t-1})}}{\sum_{j\in\Omega} \frac{\lambda_j^t - \lambda_j^{t-1}}{\ln(\lambda_j^t/b_j^t) - \ln(\lambda_j^{t-1}/b_j^{t-1})}} \tag{B.5}$$

Equation (B.5) reproduces the weights in OPI-M FBW and RW (equations 21 and 23 in the main text), with FBW's weights corresponding to the case under the assumption that $b_j^t = b_j^{t-1}$.

Equivalence to RW (Redding and Weinstein, 2020) Note that we could rewrite the change in the price index as follows:

$$\ln \hat{P}^t = \frac{1}{\sigma - 1} \ln\left(\frac{\lambda_j^t}{\lambda_j^{t-1}}\right) + \ln\left(\frac{p_j^t/(b_j^t)^{\frac{1}{\sigma-1}}}{p_j^{t-1}/(b_j^{t-1})^{\frac{1}{\sigma-1}}}\right)$$

Here, the shift in our perspective is to think of $(b_j^t)^{\frac{1}{\sigma-1}}$ as a price-equivalent taste parameter. That is, to think of the whole term $p_j^t/(b_j^t)^{\frac{1}{\sigma-1}}$ as price itself. For a generic variable x, define $\bar{x}$ as:

$$\bar{x}^t = \prod_{j\in\Omega} \left(x_j^t\right)^{\alpha_j}$$

for a set of $\{\alpha_j\}_{j\in\Omega}$ satisfying $\alpha_j > 0$ for all $j \in \Omega$ and $\sum_{j\in\Omega}\alpha_j = 1$. We follow RW and assume that $\alpha_j = 1/|\Omega|$ for all varieties j, where $|\Omega|$ is the number of varieties in the common set Ω, and that average taste remains unchanged, *i.e*, $\bar{d}_{t-1} = \bar{d}_t$. Therefore,

$$\ln \hat{P}^t = \frac{1}{\sigma - 1} \ln\left(\frac{\bar{\lambda}^t}{\bar{\lambda}^{t-1}}\right) + \ln\left(\frac{\bar{p}^t}{\bar{p}^{t-1}}\right) \tag{B.6}$$

The above expression reproduces equation (9) in RW. Our equivalent formula is based on recovering taste parameters, and plug them back as weights. First, recover demand parameters according to:

$$\left(\ln b_j^t - \ln b_j^{t-1}\right) = \ln\left[\left(\frac{\lambda_j^t}{\bar{\lambda}^t}\right)\Big/\left(\frac{\lambda_j^{t-1}}{\bar{\lambda}^{t-1}}\right)\right] - (1-\sigma)\ln\left[\left(\frac{p_j^t}{\bar{p}^t}\right)\Big/\left(\frac{p_j^{t-1}}{\bar{p}^{t-1}}\right)\right]$$

Then, replacing the above weights into equation (B.6), which reproduces equation (B.5).

Computation of Income and Own-Price Elasticity From column (4) of Table 2 (the specification with OPI-M2), we gather that $\widehat{\beta_F} = -0.0825$, $\widehat{\delta_F} = 0.0578$, and $\widehat{\delta_X} = -0.0446$. From our data, we obtain that $\overline{s_{F,h,r}^t} = 0.123$ and $\overline{\left(\ln p_{F,r}^t - \ln p_{N,r}^t\right)} = -0.0417$. Evaluating the income and price elasticities using equations (5) and (6) at the mean values, we have:

$$\frac{\partial \ln Q_{F,h,r}^t}{\partial \ln y_h^t} = 0.33, \frac{\partial \ln Q_{F,h,r}^t}{\partial \ln \tilde{p}_{F,r}^t} = -0.50. \tag{B.7}$$

Note that, in the computation of the price elasticity, the term $2\delta_X\left(\ln p_{F,r}^t - \ln p_{N,r}^t\right)$ does not contribute much, because the values of both $\widehat{\delta_F}$ and $\overline{\left(\ln p_{F,r}^t - \ln p_{N,r}^t\right)}$ are small.

C Additional Tables and Figures

Table A.3: Year Dummies for Food Engel Curve by Price Specification (Appendix)

	Dependent variable: Food Share						
Year	CPI	Constant Price	OPI-M1	OPI-M2	OPI-D	OPI-G1	OPI-G2
1996	-0.00220	0.000583	-0.000106	-0.00319	-0.00245	-0.00134	-0.00125
	(0.00181)	(0.00181)	(0.00181)	(0.00183)	(0.00181)	(0.00181)	(0.00181)
1997	-0.00390*	0.000903	-0.000916	-0.00214	-0.00421*	-0.00271	-0.00250
	(0.00188)	(0.00188)	(0.00189)	(0.00187)	(0.00188)	(0.00188)	(0.00188)
1999	-0.00784***	0.0000370	-0.00380*	-0.00211	-0.00854***	-0.00905***	-0.00851***
	(0.00185)	(0.00186)	(0.00187)	(0.00189)	(0.00185)	(0.00184)	(0.00184)
2001	-0.0102***	0.00225	-0.00237	0.00122	-0.0113***	-0.0146***	-0.0136***
	(0.00177)	(0.00179)	(0.00183)	(0.00190)	(0.00178)	(0.00177)	(0.00177)
2003	-0.00625**	0.00914***	0.00589**	0.00218	-0.00706***	-0.0102***	-0.00864***
	(0.00194)	(0.00198)	(0.00213)	(0.00201)	(0.00196)	(0.00195)	(0.00195)
2005	-0.00338	0.0165***	0.00747***	0.00459*	-0.00353	-0.00693***	-0.00475*
	(0.00193)	(0.00198)	(0.00214)	(0.00202)	(0.00194)	(0.00194)	(0.00194)
2007	-0.00687***	0.0178***	-0.00185	-0.00330	-0.00634**	-0.0119***	-0.00935***
	(0.00192)	(0.00202)	(0.00205)	(0.00194)	(0.00193)	(0.00192)	(0.00192)
2009	-0.00684***	0.0227***	0.000957	-0.000991	-0.00624**	-0.0115***	-0.00886***
	(0.00201)	(0.00210)	(0.00212)	(0.00201)	(0.00207)	(0.00201)	(0.00201)
2011	-0.00347	0.0300***	-0.00443*	-0.00219	-0.00133	-0.00599**	-0.00245
	(0.00188)	(0.00196)	(0.00189)	(0.00188)	(0.00192)	(0.00188)	(0.00188)
2013	-0.000322	0.0360***	0.00358	0.00204	0.00150	-0.00199	0.00145
	(0.00190)	(0.00204)	(0.00206)	(0.00192)	(0.00190)	(0.00190)	(0.00190)
2015	-0.00460*	0.0342***	-0.00547**	0.000417	-0.00296	-0.00679***	-0.00305
	(0.00203)	(0.00216)	(0.00199)	(0.00202)	(0.00206)	(0.00202)	(0.00202)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Base year	1995	1995	1995	1995	1995	1995	1995
Observations	23,906	23,906	23,906	23,906	23,906	23,906	23,906

Notes: Coefficients are year fixed-effect estimates from the corresponding Engel curve regressions with the same controls as in Table 2, including age, hours, and education of both head and spouse, and number of children. Robust standard errors in parentheses. Base year is 1995. Stars: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure A.2: Bias Over Time: Effect of OPI-M components

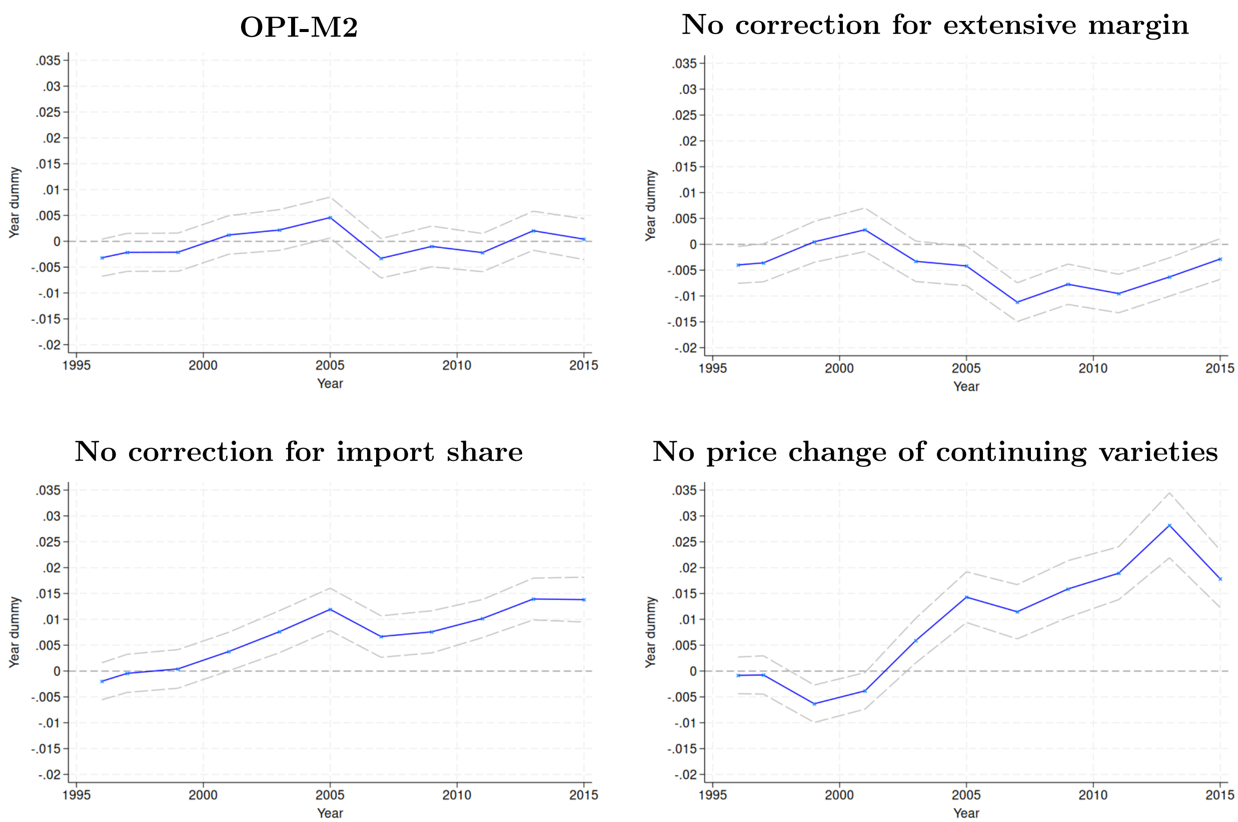

Notes: This figure presents the year dummies in OPI-M2, leaving out each of the three components one by one. These components are the correction for new and exiting varieties, changes in the import share of the economy, and observed change in prices for continuing varieties. The data period is 1995 to 2015. Year dummies are normalized to zero in 1995. We include 95% confidence intervals.

Figure A.3: Cost of Living Index increase due to 10% increase in import prices

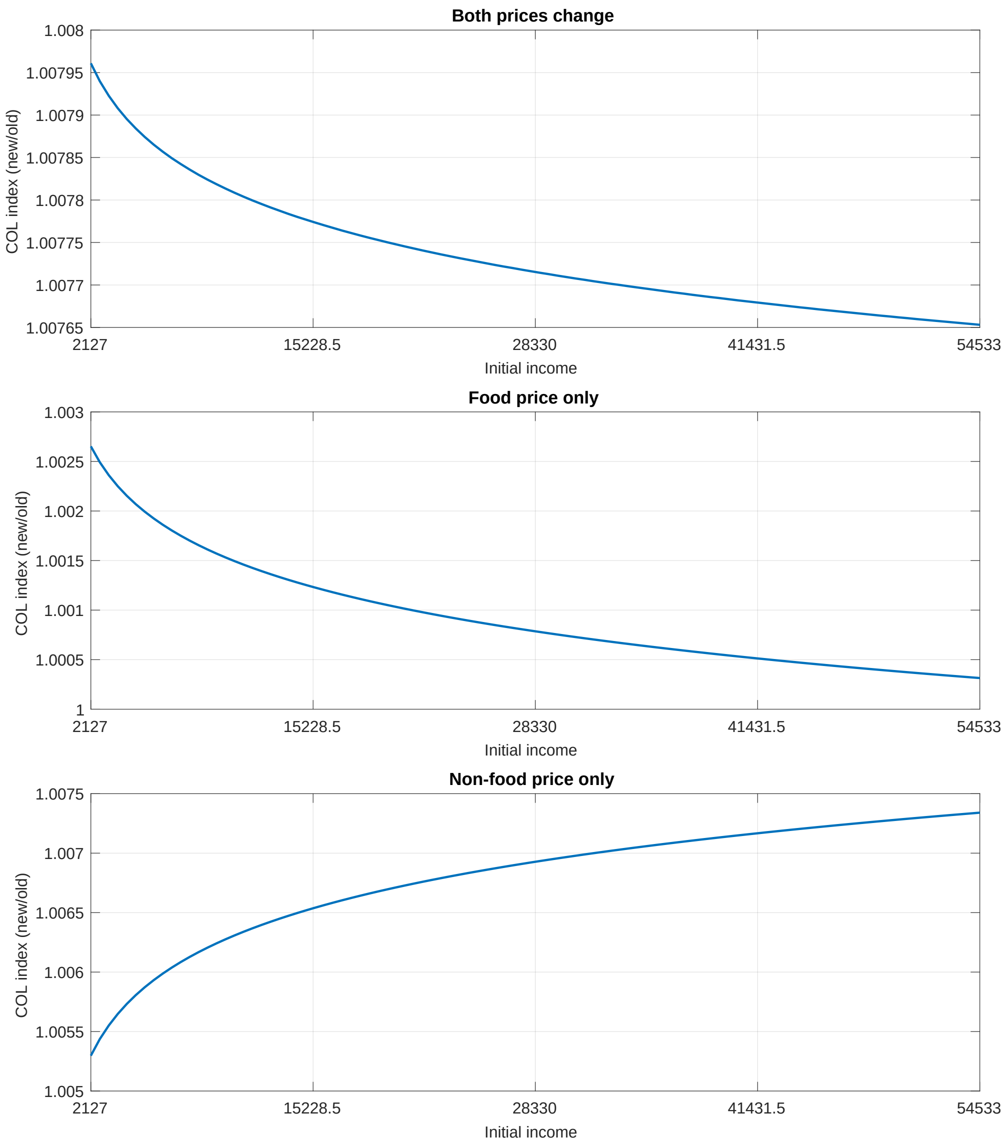

Notes: This figure shows the income required to achieve the same utility after (1) a 10% universal increase in import prices, (2) a 10% increase in food import prices, and (3) a 10% increase in non-food import prices. Since this shock generates an increase in prices for all households, the income required is higher than original income. Panels show the required income as a ratio with original income.